Azure Security Health Check

Sample report: what a 3-day read-only review finds in a small Azure estate, and what the fixes change

Prepared for: **Contoso (lab)** | Scope: subscriptions `sub-lab-prod` and `sub-lab-dmz`, 120 servers (90 Azure VMs + 30 on-premises), 38 storage accounts, 9 key vaults, 1 App Service, 1 Azure SQL database | Access used: **Reader + Security Reader** (read-only) | Date: November 2026

Prepared by: **Cuong Do**, Microsoft 365 & Azure Security Engineer (AZ-500)

LAB BUILD. This sample is assembled from the labs I publish every week (links in section 8). Every finding below was rebuilt in my own lab subscription, broken on purpose, fixed, and checked against Microsoft Learn. Screens are redrawn lab screens; names and numbers are lab values. No customer data is shown. A real report has the same structure, with your resources and your numbers.

1. At a glance

60 → 120	$1,800 → $1,000	41 → 0	58% → 64%	8	3 days
servers actually protected by Defender, out of 120 billed	Defender for Servers per month, after splitting Plan 1 / Plan 2	secrets a contractor account could read in the payments vault	secure score once the storage control turned green	findings: 2 critical, 3 high, 2 medium, 1 low	from read-only access to this report

Contoso's Azure estate runs the business, but it was switched on faster than it was secured. The three things that matter most: **half of the servers Contoso pays Defender for are not protected** (60 of 120), **a contractor account could read every secret in the payments key vault** without anyone adding it, and **the order app keeps its SQL password and storage key in plain text** where 12 people can read them. Nothing here is exotic. Each one is a default that was never changed, and each one has a fix that takes hours, not weeks.

The "after" numbers in this report are measured in my lab with the same queries I use before the fix, so the before and after can be compared line by line. That is the point of a health check: a short list, in plain English, with proof.

2. Scope and method

Item	Lab estate
Subscriptions	`sub-lab-prod` (80 VMs, storage, payments vault, order app) and `sub-lab-dmz` (10 VMs, web tier)
On-premises	30 servers in the office, including the file server `LAB-FS-01`, not yet connected to Azure Arc
Access I used	Reader and Security Reader on both subscriptions. These roles can see everything and change nothing, so the review cannot break anything.

Tools	Microsoft Defender for Cloud (the free Foundational CSPM part: recommendations, secure score, inventory), Azure Resource Graph, Azure Policy compliance details, Activity Log, and read-only `az` commands. No agents installed, no paid plan required for the review itself.
Not in scope	Penetration testing, application code review, Microsoft 365 (that is the separate Microsoft 365 Health Check), AWS or GCP
Typical real client	A 20-300 staff company, or an MSP customer, with one to three subscriptions and a few dozen servers

3. The findings, ranked by risk

#	Finding	Risk	Why it matters	Fix (effort)
1	Defender for Servers Plan 2 is billed for 120 servers but protects 60. The Defender for Endpoint integration and agentless scanning toggles are off, the plan was never enabled on `sub-lab-dmz` (a Reader clicked Save and nothing happened), 30 on-premises servers have no Azure Arc agent, and 15 DMZ extensions are stuck at *Transitioning* because egress is blocked.	Critical	$900 a month for nothing, and no EDR on the file servers and the web tier, the two places attackers like most. The auditor's "all servers have EDR" control fails.	One Owner enables the plan on both subscriptions, turns on the two toggles, Arc on the 30 office servers, open egress for the DMZ. Split Plan 2 (40 crown jewels) and Plan 1 (80 web/app). (2 days)
2	Key vault `kv-lab-pay-01` still uses the legacy *vault access policy* model. Contractor account `lab.contractor02`, given Contributor "to fix a deployment", added an access policy for itself and can read all 41 secrets, including the payment API key. 3 of 9 vaults are on this model.	Critical	In this model anyone who can write the vault can hand themselves the keys. No alarm fires, because nothing was "hacked".	Remove the policy, map each access policy to the smallest RBAC role, switch the vault to RBAC, make admins eligible in PIM, deny the old model with Azure Policy. (half a day for 3 vaults)
3	App Service `app-lab-orders` keeps the SQL password and the storage account key in plain text in its app settings. 12 people with Contributor can read them. The same password is in a dev file in Git and in a pipeline variable. Last rotated 14 months ago.	High	One leaked copy opens the orders database and every container in the storage account. Nobody rotates it because nobody knows what will break.	Managed identity on the app, one role per target, Entra-only authentication on SQL, rotate both storage keys and set `AllowSharedKeyAccess` to false. (1 day incl. test)
4	Storage account `stlabhr01` (HR documents) is "private" with a service endpoint, but its public endpoint still answers from the internet, an office IP range was added as a "fix" for VPN users, and a VM on the subnet can copy files to a storage account in another tenant.	High	A service endpoint changes the route, not the exposure. The HR share is one firewall rule away from the internet.	Private endpoint for blob, `privatelink` DNS zone, test from a VM and from the office, then disable public network access. (half a day)

5	Management ports: `vm-lab-linux-01` and 8 other VMs accept SSH 22 or RDP 3389 from Any. Just-in-time access shows *Unsupported* because there is no NSG in front of them and Plan 2 was not active.	High	Open management ports are the most common way a VM is taken over. The Activity Log already shows sign-in attempts from 40 countries.	NSG on every subnet, Plan 2 on, JIT with source "My IP" and a 1-hour window, test one request end to end. (half a day)
6	Control *Restrict unauthorized network access* scores 0 of 6 points. 38 storage accounts were "fixed" by a script, but 7 still have `networkAcls.defaultAction = Allow`; the other 31 red items were evaluation lag. Two accounts are meant to be public (static docs site) and were never exempted.	Medium	37 of 38 healthy still gives zero points. The team and the CISO spent a day arguing with a dashboard instead of fixing 7 accounts.	Set default action Deny on the 7, trigger an on-demand policy scan, exempt the 2 public accounts as *risk accepted* with an expiry date, report the score the next working day. (2 hours)
7	Alerting was never tested. Email goes only to subscription Owners (the default), only for High severity, no shared mailbox, no continuous export to the Log Analytics workspace `law-lab-mdc`, and nobody has ever created a sample alert.	Medium	A Defender plan that nobody receives alerts from is a line on the invoice, not a control.	Email to Owner + Security Admin + `soc@contoso.com` from Medium, continuous export to the workspace, sample alerts for each plan, then one narrow suppression rule for the known scanner. (2 hours)
8	The operations runbook still checks the retired recommendation *Endpoint protection should be installed on your machines* (Log Analytics agent era). The third-party EDR on `vm-lab-01` shows as unhealthy even though it works.	Low	Time spent chasing a ghost. The real check is the agentless *EDR solution should be installed* recommendation, which needs Plan 2 or Defender CSPM.	Update the runbook, read the *Discovered EDRs* column, exempt as *mitigated* where a supported third-party EDR is confirmed. (1 hour)

4. Detailed findings, with evidence

4.1 Servers and Defender coverage (finding 1)

Enabling a Defender for Servers plan is a billing decision at subscription level. Protection on each server is a separate chain: the right role to enable, the plan on the right subscription, the Defender for Endpoint integration switched on, agentless scanning switched on, and Azure Arc for anything outside Azure. Any missing link means a server that is billed but not protected. Contoso had four of the links missing at once, and none of them show on the plan page.

Figure 1. Inventory, both subscriptions plus the Arc resource group: 60 of 120 protected before, 120 of 120 after

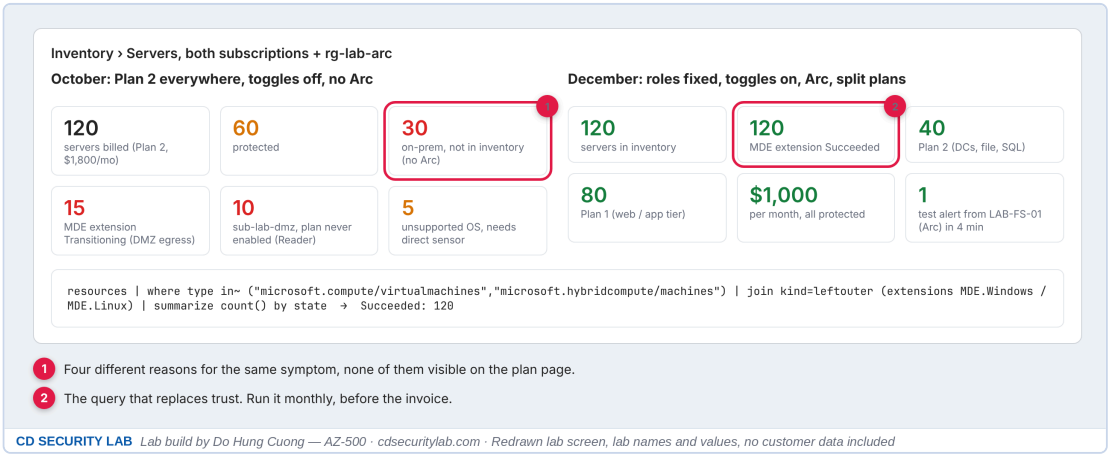

Option	Monthly	Protected
Before: Plan 2 on everything, half not onboarded	$1,800	60
Plan 2 on everything, all onboarded	$1,800	120
Plan 2 on 40 crown-jewel servers (DCs, file, SQL), Plan 1 on 80 web/app, all onboarded	**$1,000**	**120**

Plan 1 can be enabled per machine; Plan 2 can only be enabled per subscription and disabled per machine. Plan 2 also brings a pooled 500 MB per machine per day of free security log ingestion, worth up to about $2,500 a month at Sentinel prices if those machines report to the workspace. List prices as published in 2026; confirm in the pricing calculator.

How I check it (Azure Resource Graph, read-only)

```
resources
| where type in~ ("microsoft.compute/virtualmachines", "microsoft.hybridcompute/machines")
| project id, name, kind = type, subscriptionId
| join kind=leftouter (
    resources
    | where type in~ ("microsoft.compute/virtualmachines/extensions",
"microsoft.hybridcompute/machines/extensions")
    | where name in~ ("MDE.Windows", "MDE.Linux")
    | project vmId = tolower(substring(id, 0, indexof(id, "/extensions/"))), ext = name, state =
tostring(properties.provisioningState)
) on $left.id == $right.vmId
| summarize count() by state        // target after the fix: Succeeded = 120, nothing else
```

4.2 Identity and secrets (findings 2 and 3)

The payments vault. A key vault has two permission models. In the legacy *vault access policy* model the permission list is a property of the vault, so anyone who can write the vault (Contributor is enough) can add themselves. In the *Azure RBAC* model, reading a secret is a normal role assignment that only Owner or User Access Administrator can grant. Same role, two very different meanings.

Figure 2. Before: the contractor added Get and List for itself. 41 secrets readable

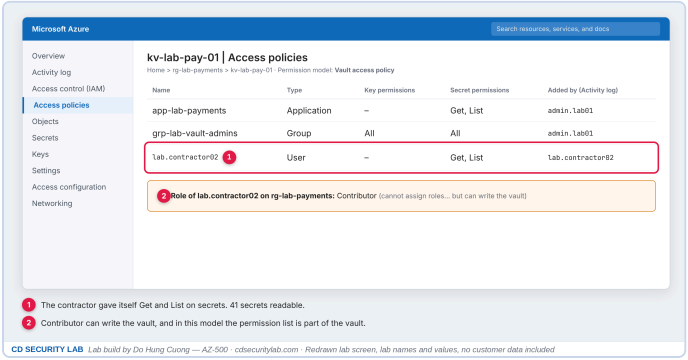

Figure 3. After: RBAC model, one smallest role per app, admins eligible in PIM, policy set to Deny

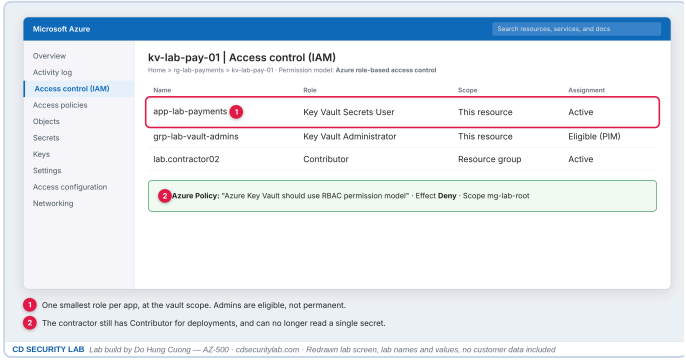

Find every vault still on access policies (all subscriptions you can see)

```
resources
| where type == "microsoft.keyvault/vaults"
| extend rbac = tobool(properties.enableRbacAuthorization)
| where isnull(rbac) or rbac == false
| project name, resourceGroup, subscriptionId, policies = array_length(properties.accessPolicies)
// lab: 3 of 9
```

Fix order matters: list who is in each policy, give the same access with RBAC roles first, then `az keyvault update --name kv-lab-pay-01 --enable-rbac-authorization true`. The moment you switch, all access policies stop working. Then assign the built-in policy "Azure Key Vault should use RBAC permission model", Audit first, then Deny.

The order app. Moving the SQL password into a vault would help, but the better question is why the app needs a password at all. A managed identity is an identity Azure creates for the app; the platform proves who the app is, and there is nothing to store, copy or rotate.

Figure 4. Before: SQL password and storage key in app settings, readable by 12 Contributors, last rotated 14 months ago

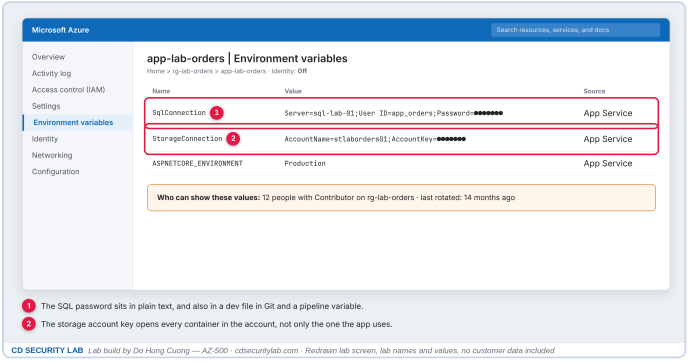

Figure 5. After: identity on, one role per target, no secret left in the app, old password and keys closed

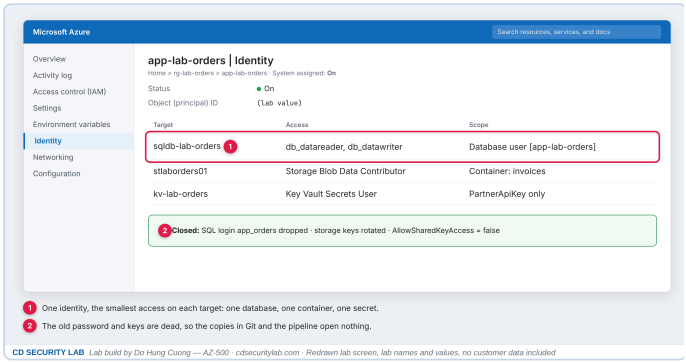


```
az webapp identity assign -g rg-lab-orders -n app-lab-orders
# SQL: CREATE USER [app-lab-orders] FROM EXTERNAL PROVIDER; ALTER ROLE db_datareader ADD MEMBER
[app-lab-orders]; ...
# Storage: Storage Blob Data Contributor on the ONE container the app uses, not the whole account
# Then close the old doors: Entra-only auth on the SQL server, rotate both keys,
AllowSharedKeyAccess = false
```

4.3 Network exposure (findings 4 and 5)

The HR storage account. A service endpoint says "traffic from this subnet to Azure Storage goes over the Azure backbone". It does not close the public endpoint, it does not work for the office over VPN, and it does not care which storage account the traffic goes to. A private endpoint gives *this one* account a private IP inside the network, and only then can public access be switched off.

Figure 6. Before: service endpoint on, public endpoint still answering, office IP range added as a "fix"

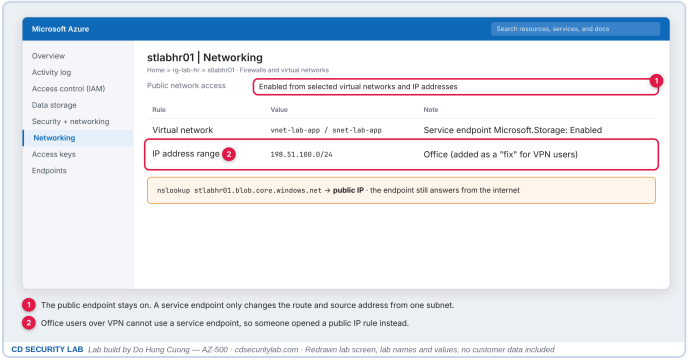

Figure 7. After: private endpoint, private DNS zone, public network access disabled

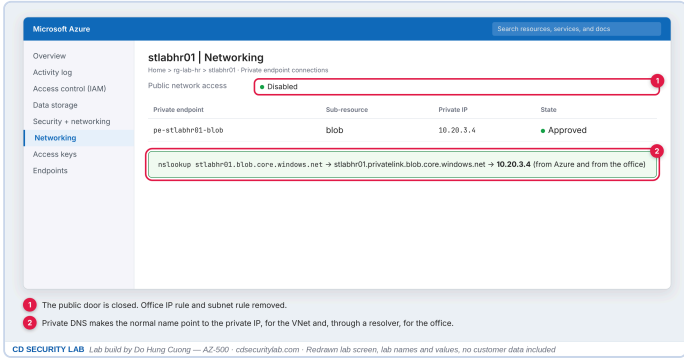


```
az network private-endpoint create -g rg-lab-hr -n pe-stlabhr01-blob --vnet-name vnet-lab-app --
subnet snet-lab-pe \
  --private-connection-resource-id "$(az storage account show -n stlabhr01 --query id -o tsv)" --
group-id blob --connection-name stlabhr01-blob
nslookup stlabhr01.blob.core.windows.net     # expected: alias to
...privatelink.blob.core.windows.net, private IP
az storage account update -g rg-lab-hr -n stlabhr01 --public-network-access Disabled   # only after
both tests pass
```

Management ports. Just-in-time VM access keeps SSH and RDP closed in the NSG and opens them for one admin, from one IP, for a short time, on request. It needs Defender for Servers Plan 2 and an NSG or Azure Firewall in front of the VM; without those the portal shows *Unsupported*, which is what Contoso saw and gave up on.

Figure 8. After: JIT request on vm-lab-linux-01, port 22, source My IP, 1 hour, logged in the Activity Log

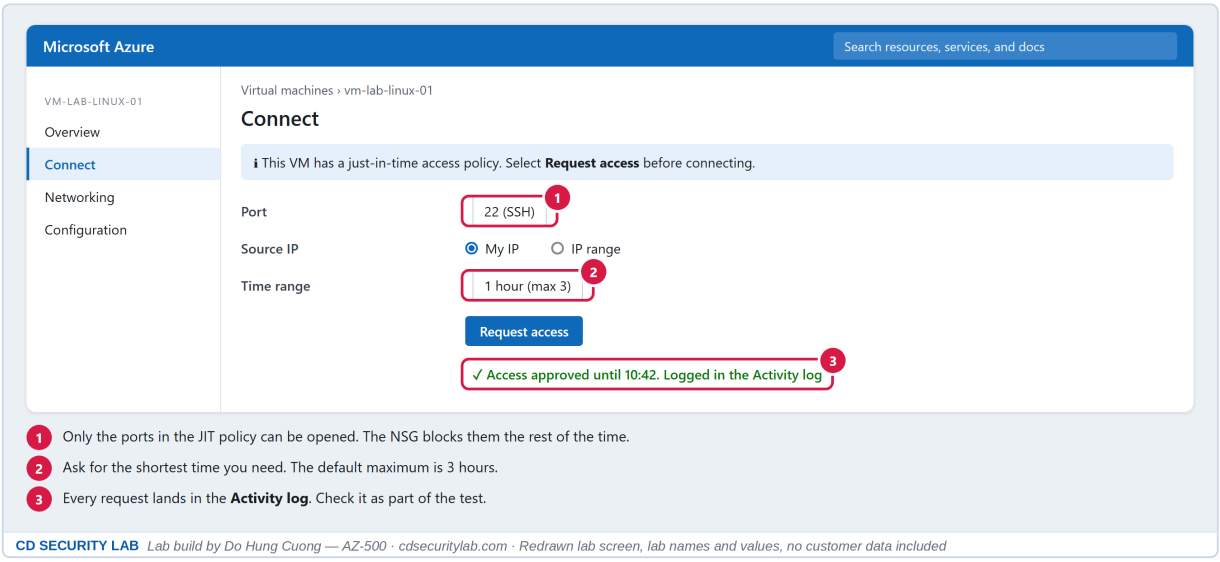

Defaults when JIT is enabled from the VM page: RDP 3389 or SSH 22, 3 hours maximum, source *Any*. Change "Any" to the admin IP range. Test it like a user: request, connect, then find the request in the Activity Log.

4.4 Posture and secure score (finding 6)

A recommendation is not a live view. Azure Policy re-evaluates a changed resource in about 15 minutes and everything else every 24 hours; recommendations refresh on their own interval; secure score controls are recalculated every 8 hours; an exemption takes up to 24 hours. And a control gives points only when **all** its resources are healthy. Contoso fixed 38 storage accounts on Thursday at 17:00; the Friday 09:00 report still showed 31 unhealthy and the same score. 24 of those were lag. 7 were real: the script set `publicNetworkAccess = Disabled` but left the firewall default on Allow.

Figure 9. The four clocks between a fix and a green tick, and why Friday's report could not be green yet

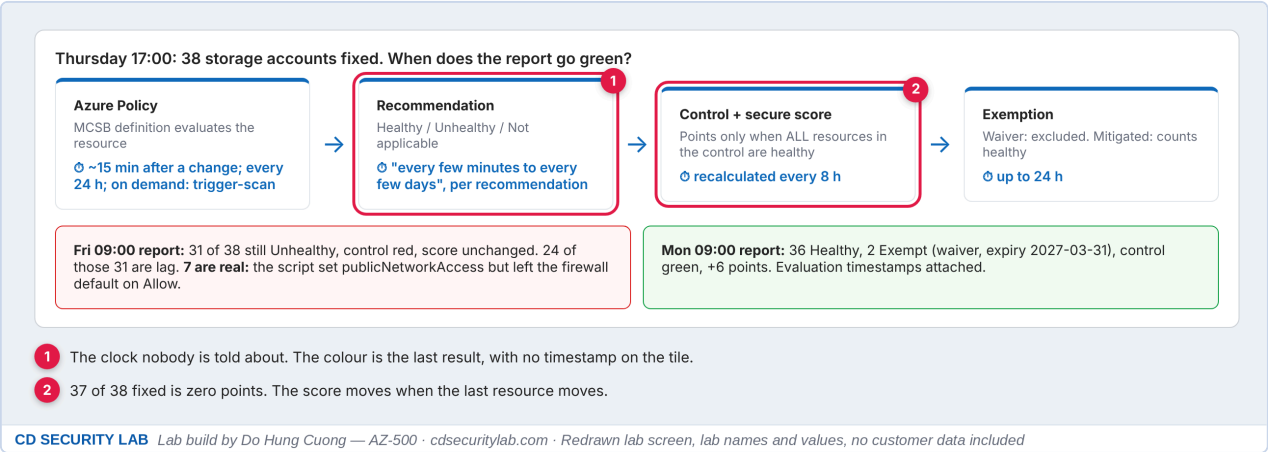

1. The clock nobody is told about. The colour is the last result, with no timestamp on the tile.
2. 37 of 38 fixed is zero points. The score moves when the last resource moves.

CD SECURITY LAB *Lab build by Do Hung Cuong — AZ-500 · cdsecuritylab.com · Redrawn lab screen, lab names and values, no customer data included*

Figure 10. Same control, Friday vs Monday: 31 unhealthy and 0/6 points, then 36 healthy, 2 exempted, 6/6 points, score 58% to 64%

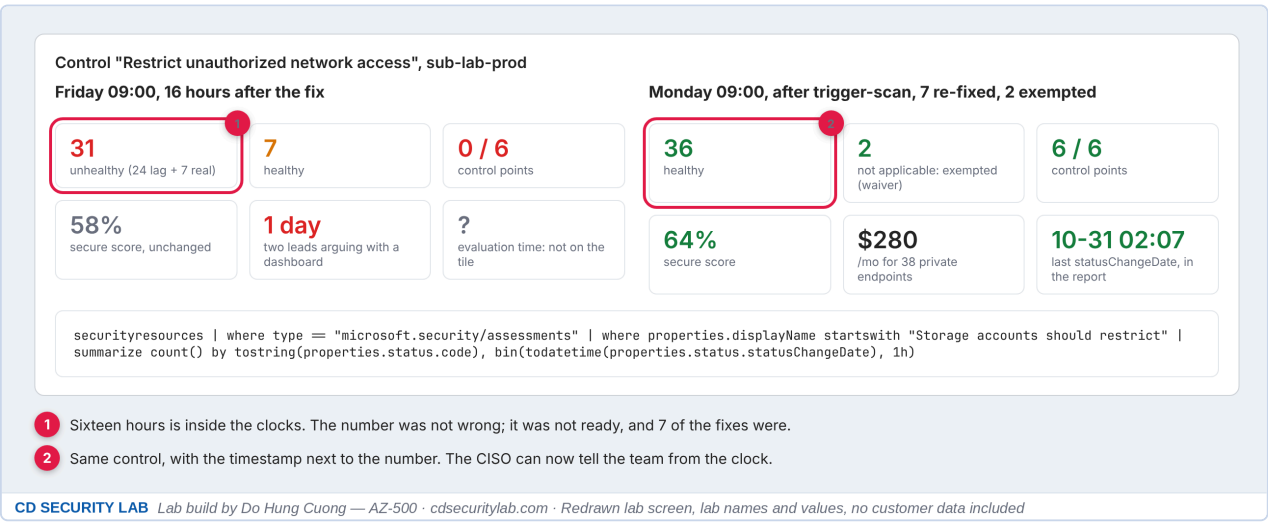

1. Sixteen hours is inside the clocks. The number was not wrong; it was not ready, and 7 of the fixes were.
2. Same control, with the timestamp next to the number. The CISO can now tell the team from the clock.

CD SECURITY LAB *Lab build by Do Hung Cuong — AZ-500 · cdsecuritylab.com · Redrawn lab screen, lab names and values, no customer data included*

```
for sa in $(az storage account list -g rg-lab-storage --query "[?
networkRuleSet.defaultAction=='Allow'].name" -o tsv); do
  az storage account update -g rg-lab-storage -n "$sa" --default-action Deny; done
az policy state trigger-scan --resource-group rg-lab-storage --no-wait      # do not wait for the
24-hour cycle
```

Figure 11. The two public accounts: exemption as "risk accepted" with an expiry date, not "disable the recommendation"

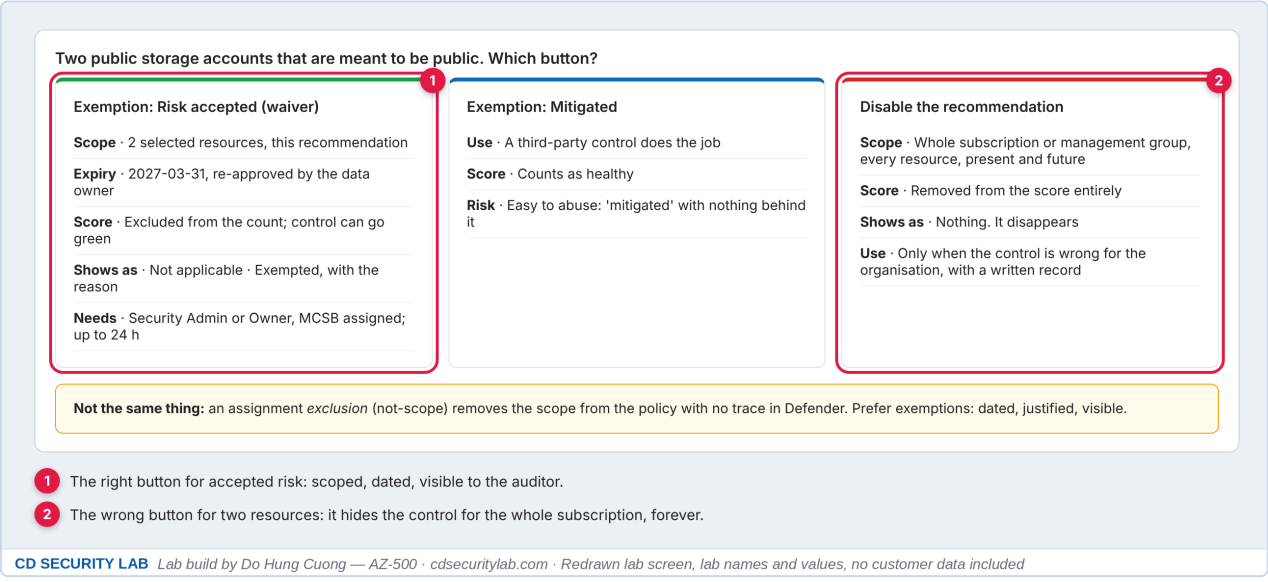

1. The right button for accepted risk: scoped, dated, visible to the auditor.
2. The wrong button for two resources: it hides the control for the whole subscription, forever.

CD SECURITY LAB *Lab build by Do Hung Cuong — AZ-500 · cdsecuritylab.com · Redrawn lab screen, lab names and values, no customer data included*

My rule: every exemption has a reason, an owner and an expiration date. An exemption with no end date becomes the new normal. Disabling a recommendation removes it from the score for every resource, present and future, and leaves no trace.

4.5 Alerting and notifications (finding 7)

Turning on a Defender plan does not prove anyone will see an alert. By default only subscription Owners get an email, only for High severity, and email is rate-limited (4 High alerts per address per day). The test is simple: create sample alerts and check that they arrive everywhere alerts should go: email, continuous export, and the SIEM.

Figure 12. After: roles plus a shared mailbox, alerts from Medium, attack paths from High

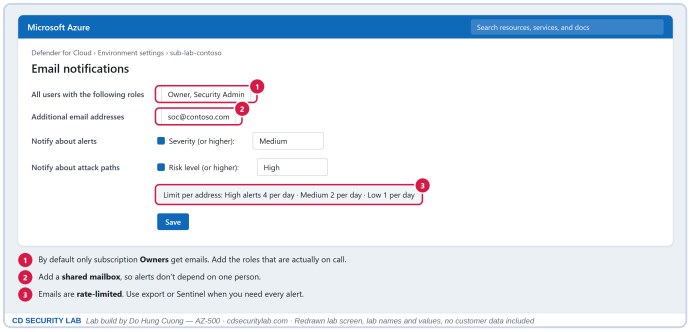

Figure 13. Sample alerts for each enabled plan, the end-to-end test of the alert path

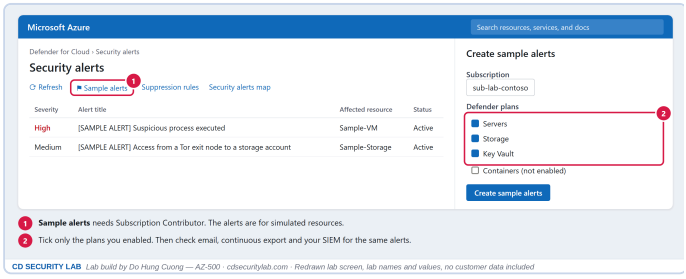

If the sample alert is in the portal but not in the workspace or Sentinel, the export or connector is broken. Better to find out now than during an incident. Suppression rules come after: one alert type, narrowed to one host or IP, with an expiry, simulated before applied.

4.6 Retired checks in the runbook (finding 8)

The recommendations *Endpoint protection should be installed on machines* and *Endpoint protection health issues should be resolved* relied on the Log Analytics agent, which retired in November 2024. Today the check is agentless: *EDR solution should be installed on Virtual Machines* and *EDR configuration issues should be resolved*, and it needs Defender for Servers Plan 2 or Defender CSPM. Open the recommendation, tab *Healthy resources*, column *Discovered EDRs*: if the third-party EDR is listed, it is fine. If it is real but unsupported, exempt the machine as *resolved through third-party service (mitigated)*, which counts as healthy.

5. Recommended plan

When	Actions	Owner
Week 1, quick wins (2 days)	Findings 1, 2, 6, 7: one Owner enables Defender for Servers on both subscriptions and switches on the two integration toggles; remove the contractor's access policy and switch the 3 vaults to RBAC; set default action Deny on the 7 storage accounts and trigger a policy scan; email notifications to roles plus `soc@contoso.com`, continuous export, sample alerts	Cuong Do with Contoso IT lead (Owner)
Week 2 (3 days)	Findings 3, 4, 5: managed identity for the order app, Entra-only SQL auth, key rotation and shared key off; private endpoint and private DNS for `stlabhr01`, then public access off; NSGs, Plan 2 and JIT on the 9 exposed VMs, one request tested end to end	Cuong Do with the app owner and the network contact

Weeks 3-4	Finding 1 (rest) and 8: Arc agent on the 30 office servers, DMZ egress to the Defender for Endpoint endpoints, split Plan 1 / Plan 2 by server class, update the runbook, exempt the confirmed third-party EDR as mitigated; Azure Policy in Audit for Key Vault RBAC, private link on storage, managed identity on App Service, then Deny	Contoso IT with my runbooks
Day 30	Re-run every query in section 6, compare with this report line by line, report the secure score with its evaluation timestamp	Cuong Do

Cost of the fixes in the lab: private endpoints at about $7.30 a month each, Defender for Servers from $1,800 down to $1,000 a month after the split. Everything else is configuration only.

6. How the result is proven

Every "after" claim in this report is backed by a query that anyone with Reader can run again. The four I use most:

```
// Which subscriptions actually have the Servers plan
securityresources | where type == "microsoft.security/pricings" | where name == "VirtualMachines"
| project subscriptionId, tier = tostring(properties.pricingTier), subPlan =
tostring(properties.subPlan)

// Secure score per subscription
SecurityResources | where type == 'microsoft.security/securescores'
| extend current = properties.score.current, max = todouble(properties.score.max)
| project subscriptionId, current, max, percentage = ((current / max)*100)

// Storage network control: status with the timestamp next to the number
securityresources | where type == "microsoft.security/assessments"
| where properties.displayName startswith "Storage accounts should restrict network access"
| extend status = tostring(properties.status.code), changed =
todatetime(properties.status.statusChangeDate)
| summarize count() by status, bin(changed, 1h) | order by changed desc

// Any access-policy write on a vault after the switch is a finding (Sentinel, Azure Activity
connector)
AzureActivity | where OperationNameValue in~ ("MICROSOFT.KEYVAULT/VAULTS/ACCESSPOLICIES/WRITE",
"MICROSOFT.KEYVAULT/VAULTS/WRITE")
| project TimeGenerated, Caller, ResourceGroup, _ResourceId, ActivityStatusValue | order by
TimeGenerated desc
```

7. What you get

- This report (PDF) with findings ranked by risk, in plain English, each with the evidence and the exact fix.
- A 30-day fix plan with owners and effort, and the runbooks for every fix so your team can run them.
- The evidence pack: Resource Graph queries, RBAC and NSG exports, Defender recommendations export, shared through a link that expires after 30 days.
- A 30-minute walkthrough call, and a re-check on day 30 with the before/after comparison.

8. Sources: the labs behind this sample

Each finding above is a lab I built, broke and fixed, with the write-up published on GitLab (and the short version on Dev.to and LinkedIn). All checked against Microsoft Learn in October and November 2026.

Finding	Lab write-up
1, 8	You switched on Defender for Cloud. Half the servers never showed up. and Defender for Cloud in one page (plus a mock case)
2	SC-500 deep dive #1: your Contributors can read your Key Vault secrets
3	SC-500 deep dive #2: the best secret is the one you never had
4	SC-500 deep dive #3: "private" storage that still had a public door
5, 7	A Defender for Cloud readiness checklist: what to test before you call it done
6	You fixed it yesterday. Defender for Cloud still says it's broken.

Sample report built from my published lab write-ups. Contoso, resource names and figures are lab values; no customer data.
Cuong Do · AZ-500 · Microsoft 365 & Azure Security · cdsecuritylab.com · gitlab.com/cuongdo.hcm/azure-security-weekly · upwork.com/freelancers/cdsecuritylab