REPORT · FIRST EDITION

The Hidden Azure Work Index

How client environments compare, what's open, and what good looks like.

Every Azure environment holds work that needs attention. This report shows where the biggest areas of that work sit, how client environments compare, and where better cloud operations can deliver more value to your customers and your P&L.

85.6%

of the open work is security and reliability

417

client subscriptions with open findings, in 66 ANZ companies

spotto

CONTENTS

What’s in this report.

START HERE

The short version, how each role in your practice uses it, and what the numbers mean.

- Foreword 3
- Highlights 4
- How each role in your practice uses this report 6
- How to read the numbers 7

1 · THE BENCHMARK

How client environments stack up, and what the best-kept quartile has done.

- Where the open Azure work sits 9
- What a client’s Azure spend tells you 10
- How big the work is 11
- What good looks like 12

2 · THE WORK

Two kinds of work, six worked examples and the five bundles.

- Two kinds of work 15
- Worked examples 16
- The five service bundles 17

3 · PUTTING IT TO WORK

Six ways to use what you find, what good management gives back, and how to run this on your own clients.

- Six ways to put this to work 19
- What good management gives back 20
- Key takeaways 21
- Run this on your own clients 22

APPENDIX

How much clients differ, and where the numbers come from.

- How much clients differ 24
- Method and scope 25



What does good look like?

When the RMM arrived, most of the field work went with it. The best MSPs put the time saved into broader scope: networking, then cloud, then security and compliance.

AI is the next wave, and the scope it opens up is outcomes. An outcome only counts if your client can see it, and seeing it starts with a baseline.

In Azure, that baseline has been missing. A typical mid-spend subscription in this report has 273 open findings, and Advisor score and Secure Score rate each environment against Microsoft's checks. None of that tells you how a client compares with similar environments at other MSPs.

With that many findings on screen, it's easy to feel behind... right?

Not quite. We were curious, so we looked: what does a well-managed Azure environment look like across ANZ?

We compared hundreds of client subscriptions, each with a provider looking after it, counting high-impact findings for every 100 resources so big and small environments compare fairly.

A well-managed environment still has work open. It just has far less: about a third as many high-impact findings, for its size, as the environments with the most. Most of the open work is about keeping clients secure and running: 85.6% is security and reliability, and 3.5% is cost. The biggest gaps between the best-kept and the rest are apps that keep running when one server goes down, storage that survives an outage, and backups.

That's your baseline. Show a client where they sit, agree where they'd like to be, and report the progress each month. Your team knows what to tackle first, and your client can see it moving. With Spotto, your team sees the next job in each client and puts it to them with the why, the effort and the words drafted.

Read this against a few of your own clients. Where would they sit?

The Spotto team

The five findings that matter most.

We looked at 417 client subscriptions across 66 ANZ companies to see what work was still open in their Azure environments. Here are five things that stood out, and what they could mean for your practice.

85.6%

of open findings are security and reliability.

Just 3.5% are related to cost.

Most of the work keeps clients safe

Cost is often the easiest conversation to start with. There's a clear number attached to it, and clients understand the benefit. But there's a lot more work to discuss, from keeping systems secure to making sure they keep running when clients need them.

Pro tip: Agree upfront that any cost savings, or a percentage of the cost savings, will be allocated to other value-added activities like reducing risk.

~3×

as many findings per dollar in the lowest-spend clients as in the highest-spend

Small clients can hold more work than you'd expect

Clients that spend more on Azure have more work overall. For every dollar they spend, though, the lowest-spend clients carry close to three times as many findings, and around 40% of findings are high-impact at every size.

Don't judge the work available by the Azure bill alone. Smaller clients can still hold high-impact jobs worth scoping.

Pro tip: Frame the work around the outcomes and the client's business goals, which makes approval easier. Much of the work repeats across environments, so you can tackle the same fix for several clients at once, or package it as a service, such as a security month offered to every client.

69%

of high-impact findings are medium effort

Most high-impact jobs are a manageable size

They're bigger than a quick fix, but smaller than a major project.

Many of these jobs can be scoped, priced and quoted individually.

Pro tip: Group related jobs into a service to make them easier to plan.

5

bundles the recurring jobs fall into

The same jobs come up for most clients

They fall into five bundles: identity and access, network security, resilience and recovery, database posture, and cost and cleanup.

Scope and price each bundle once, then reuse it across your client base. That saves time on scoping new work and can help your margin.

Pro tip: Bundle smaller jobs together to put unused retainer or block hours to work.

70% fewer

open high-impact findings in the best-kept environments, for the same number of resources

The best-kept environments have about a third as much open work

A typical environment in the best-kept quartile has 67 high-impact findings per 100 resources, compared with 227 in the least-kept quartile. That's about a third as many findings for the same number of resources.

You can show clients where their Azure environment sits compared with others. For example, 69% of subscriptions in the best-kept quartile have backup configured, compared with 41% in the least-kept.

Pro tip: Track progress with each client over time, and agree a realistic target together. The environments making the most progress cut their open high-impact work by about a fifth in three months.

> **One limit to know** Every client is different. Some of what is open could or should be done, and some is open for a reason the client has chosen.

Each person in your practice uses this report differently.

Whether you're running the business, managing customers or leading the technical team, here's what to look for in the report and what it means for you.

BUSINESS OWNER

If you run the practice, look at how much of the work repeats across clients, and how that can become a service you scope once and reuse. You can also compare a new client's environment with the best-kept quartile.

START ON Two kinds of work

PRACTICE LEAD

Most high-impact findings take a medium amount of effort to fix. Check what each client's agreement already covers, then scope and bundle the remaining work into a roadmap. Smaller jobs can also help fill unused retainer or block hours.

START ON How big the work is

SERVICE DELIVERY LEAD

Track each client's high-impact findings per 100 resources month to month, and use that trend in service reviews. Check what has changed before treating a rising count as new work. The environments making the most progress cut open high-impact work by about a fifth in three months.

START ON What good looks like

ENGINEER

The same fixes come up across client environments, so document and script repeatable work where it makes sense. Check whether a setting is a deliberate choice before changing it. Spotto also estimates the effort for most fixes.

START ON Two kinds of work

PRE-SALES

Compare a prospect's environment with others in the report to see where they sit. Most subscriptions have at least one quick win, which gives you a useful starting point, and the gap to the best-kept quartile helps shape an initial proposal around the prospect's needs.

START ON How much clients differ

What the words and numbers mean.

THE WORDS

Finding

One thing Spotto flags as worth fixing on a resource.

Counted per resource

One check that fails on 400 resources counts as 400 findings. Where a page counts each recommendation once instead, it says so.

Subscription, client company

A subscription is one Azure subscription. A client company can have several: about six each here.

Impact

Spotto's rating of how much a finding matters: high, medium or low.

Effort

Spotto's rating of how much time the fix takes: low, medium or high.

Quick win

A high-impact finding with a low effort rating.

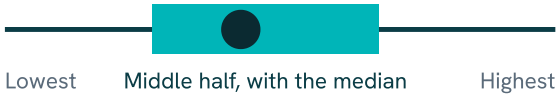
Policy check

A compliance finding from a policy assigned in Azure: a regulatory standard, Microsoft's security baseline, or the MSP's own rules. Each carries one blanket rating, so impact and effort figures leave them out.

Bundle

A group of findings you scope and sell together. There are five, set out in 'The five service bundles'.

THE NUMBERS



Median

Line everything up in order, and the median is the one in the middle. Half sit above it and half below, so one very large environment can't drag it the way it drags an average.

Middle half

The range the middle 50% fall in, between the lower and upper quartile marks. It leaves out the extremes at each end.

Best-kept quartile

Line every environment up by open high-impact findings per 100 resources, then split the line into four. The best-kept quartile carries the fewest, the least-kept quartile the most. See What good looks like.

Thirds

Subscriptions split into three equal groups by monthly Azure spend, in 'What a client's Azure spend tells you'.

Per 100 resources

A count scaled to the size of the environment, so big and small environments compare fairly.

Job size

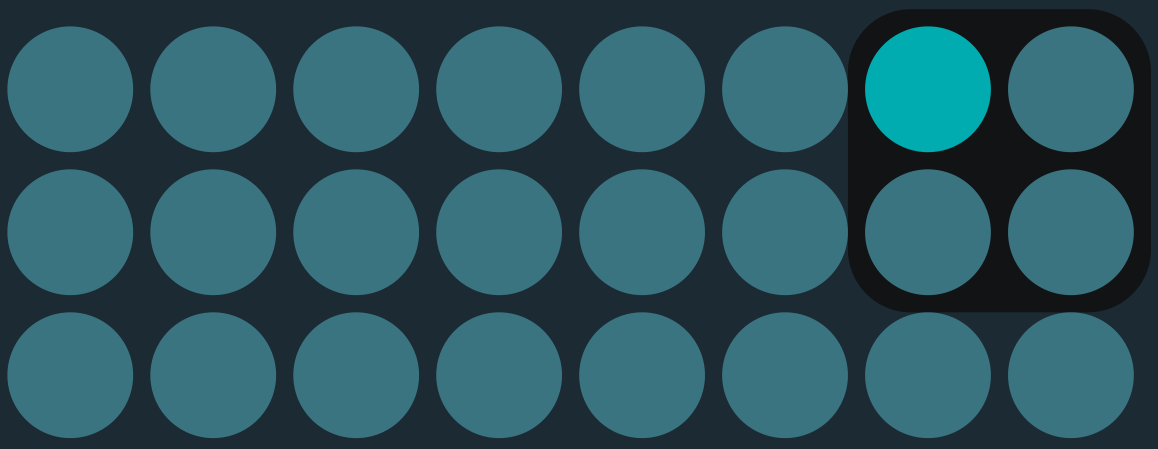
Spotto's effort rating shown as a size: small is low effort, medium is medium, large is high effort. A size assumes nothing unexpected depends on the setting.

Everything is read from Spotto on 15 September 2026: 417 client subscriptions with open findings, in 66 ANZ client companies. Each currency is kept separate. Method and scope, at the back, says where each figure comes from.

1 THE BENCHMARK

How client environments compare.

What's open across 417 client subscriptions, and what the best-kept environments have done.



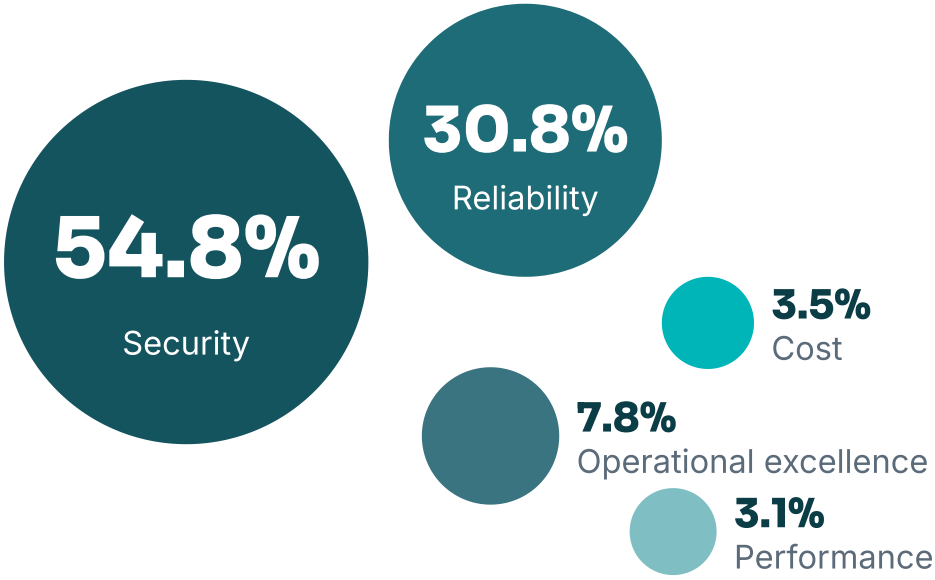
WHAT'S OPEN

Where the work sits, what spend tells you, and how big the work is.

WHAT GOOD LOOKS LIKE

How the best-kept quartile compares, and what changed over three months.

More than half of open Azure work is security work.



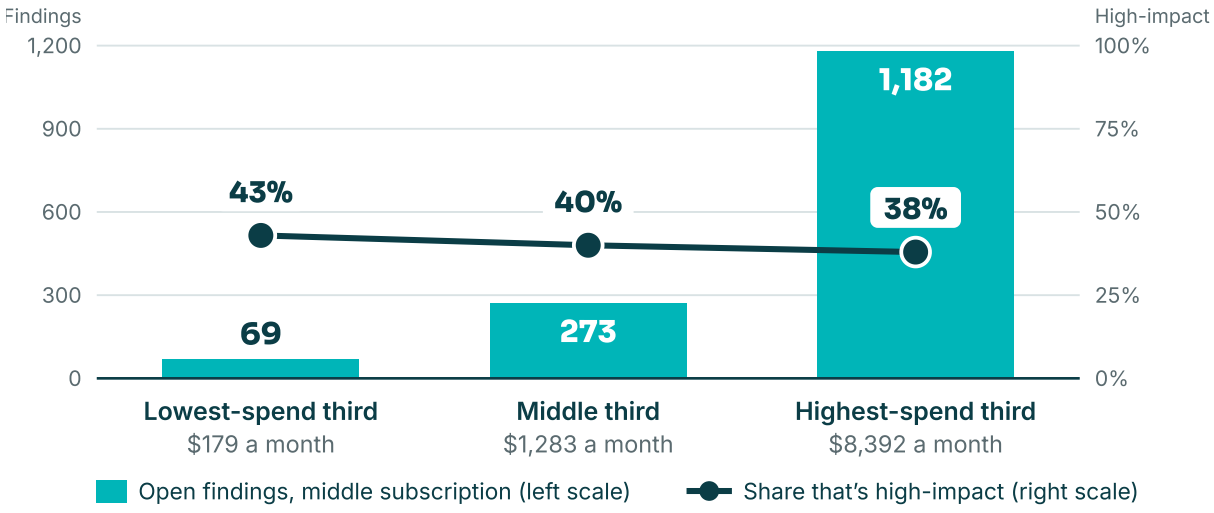
417 client subscriptions, by type of work · circles sized by share

Why there's no compliance circle. Azure can flag the same issue twice: once as a security or reliability finding, and again as a compliance policy check. We counted each issue once, under the type of work it relates to. Tagging and logging checks sit under operational excellence.

What this means for you

Security and reliability account for most of the open work across these environments. Look for the same types of findings across your client base. When the work keeps coming up, you can scope it, price it and deliver it as a repeatable service.

Almost every environment carries high-impact work, whatever the bill.



Does a bigger bill mean more work?

Yes. Bigger clients generally have more findings, but smaller clients can still have plenty of work for what they spend. The lowest-spend clients have close to three times as many findings per dollar as the highest-spend clients.

Does it mean more high-impact work?

Bigger clients have more high-impact findings overall, but the share is similar across all client sizes. Around four in ten findings are high-impact, whether a client is in the lowest-spend group or the highest.

The amount of work also varies widely between clients. For every $1,000 of monthly Azure spend, clients average about 200 findings. A few clients carry a lot of work for their bill, and they pull that average up, so the median gives a fairer picture: the client in the middle has about 150. Half of clients have between about 90 and 250. Two clients with similar bills can have very different amounts of work to address.

HOW BIG THE WORK IS

Most high-impact jobs take a medium amount of effort.

	LOW EFFORT	MEDIUM EFFORT	HIGH EFFORT
High impact	2,416	**8,514**	1,375
Medium impact	3,509	8,603	247
Low impact	2,413	936	0

How much effort does high-impact work take?

Of the 12,305 high-impact findings rated, 69% are medium effort, 20% are low effort and 11% are high effort.

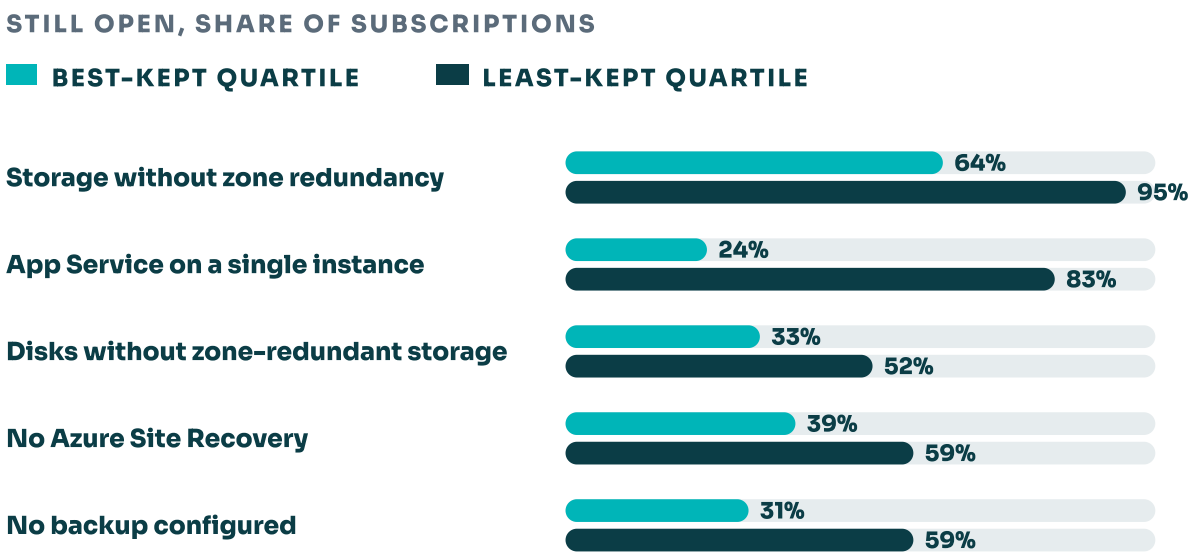
What does medium effort look like?

These jobs take more than changing a setting, but less than a migration. They need planning and engineering time, but many can be scoped and quoted individually.

What this means for you

Look for related jobs across your clients and group them into repeatable services. Smaller fixes can also help fill unused retainer or block hours, while larger jobs may suit a project.

The best-kept environments are much further through the basics.



What the best-kept quartile is

We ranked environments by high-impact findings per 100 resources, so big and small environments compare fairly. The best-kept quartile has 91 or fewer. The least-kept has 171 or more, almost twice as many for the same number of resources.

What it shows

The gap is widest in storage resilience, app availability and recovery. The best-kept environments still have work open, just far less of it.

How it helps you

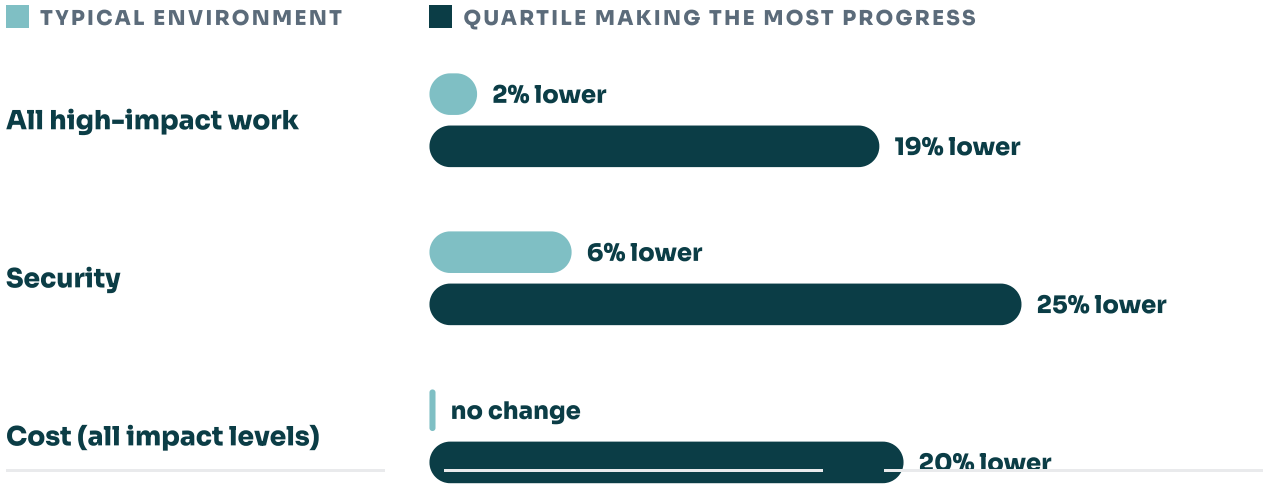
Compare each client's environment with these benchmarks to see whether it sits closer to the best-kept or the least-kept quartile.

Use that as a compelling reason to help the client move forward with change.

Try this with one client. Count their high-impact findings per 100 resources and see which quartile they sit in. Then use the table to identify where their environment differs from the best-kept group. Check whether each subscription is a production subscription before recommending changes: Site Recovery and backup add to the bill, dev or test environments may leave them off on purpose, and zone redundancy needs a region with availability zones (Australia Southeast has none).

The environments making the most progress cut open high-impact work by about a fifth in three months.

CHANGE IN OPEN FINDINGS OVER 90 DAYS



What we measured

We tracked open high-impact findings across 201 Azure subscriptions over the last 90 days, excluding days affected by changes to Spotto's checks and environments where the number of resources fell by 20% or more.

What it shows

The typical environment stayed roughly level over the 90 days. The quartile making the most progress reduced its open high-impact findings by 19%.

How it helps you

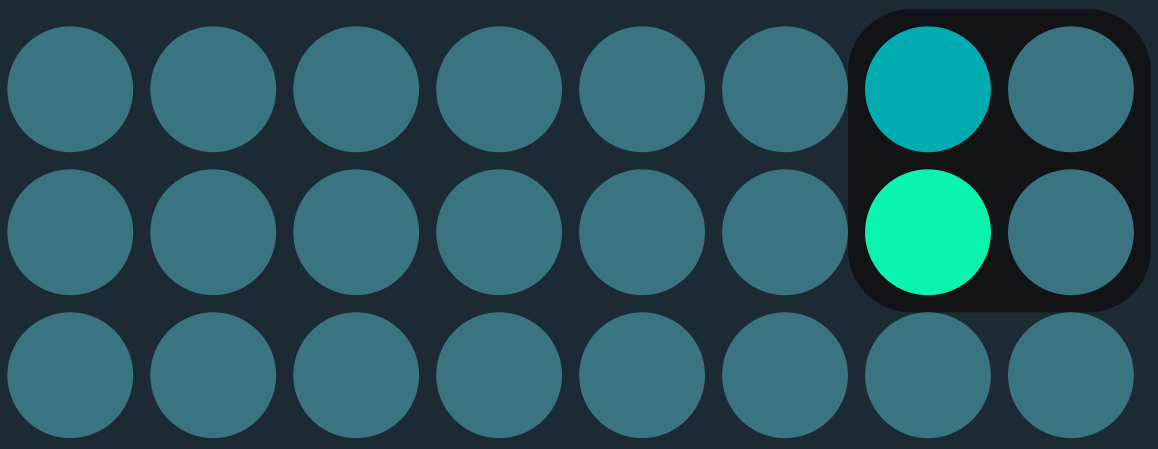
The quartile making the most progress gives you a reference point for what progress can look like. Use it in monthly client reviews to show how open findings have changed and agree what to tackle next. Every client starts from a different place, so read the trend alongside their priorities.

Use it with a client. Compare their open findings over time. Talk through what's changed, what work has been addressed, what has arrived since and where to focus next.

2 THE WORK

What the work looks like.

Some work repeats across environments, and some needs its own scope. Here are both, with six worked examples and the five bundles the recurring work falls into.



TWO KINDS OF WORK

Recurring findings, and individual jobs you can scope and price.

WORKED EXAMPLES

Six jobs, with the fix, the hours and the risk.

THE FIVE BUNDLES

Recurring work grouped into services you scope once.

TWO KINDS OF WORK

Some work repeats across environments. Other jobs need their own scope.

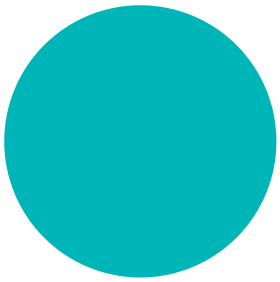


ROUTINE

Recurring findings you can bundle

These are the high-volume findings that keep turning up across client environments. A virtual machine with no backup, a disk without zone-redundant storage, or an app without a managed identity. Azure Advisor or Policy flags each issue, and your engineers can use a consistent approach to address it.

Group related findings into bundles you can scope once and reuse across clients. The five bundles follow the worked examples.



SPECIFIC

Individual jobs you can scope and price

Some findings need a closer look at a particular resource, its dependencies and what the client needs. A log workspace costing more than it should, or an App Service plan running on an older, more expensive generation, for example.

Each gives you a specific job to investigate, price and put to the client, with a clear outcome to show.

The next page works through six examples and what's involved in each.

HOW MUCH OF THE WORK REPEATS

42

There are 368 different kinds of high-impact finding, but 42 of them make up three-quarters of the work. Each of those 42 turns up at half the client companies or more, so a fix you script once gets used again and again.

High-impact findings across 66 client companies, 15 September 2026, counted by resource; policy checks excluded.

High-value work is a job you can scope.

Make storage zone-redundant

‘Every copy of the data sits in one data centre.’

The account keeps all its copies in a single data centre, so a zone outage can take the data offline. Check what the account holds and how critical it is, confirm the region supports zone-redundant storage, then change the redundancy setting and note the cost change. Some account types can’t be converted in place and need a migration.

About 8 hours · the storage bill goes up

Turn on HTTP/2 for web apps

‘Your web app is making requests wait in line.’

With HTTP/2 off, browsers fall back to HTTP/1.1, which can slow pages that load many assets. The gain depends on the app and on what sits in front of it. Turn it on in a staging slot, compare latency and errors, then roll it out.

About 3 hours · low risk

Put Log Analytics tables on the right plan

‘A premium filing cabinet for papers you rarely open.’

Some high-volume tables sit on the full Analytics plan, or are kept longer than needed. Rank tables by cost, check which alerts, dashboards and Sentinel rules use them, and move eligible ones to Basic Logs, which can’t drive Sentinel analytics rules. Before shortening retention, check compliance needs or move data to long-term retention.

About 8 hours · medium risk: check alerts and saved queries

Switch apps and services to managed identities

‘Passwords in the config, where an identity would do.’

The resource signs in to other services with secrets or keys that someone has to store and rotate. Turn on a managed identity, give it only the roles it needs, update the app to use it, then remove the old secrets.

About 12 hours · test each connection after the switch

Add a second App Service instance

‘One server, so routine maintenance means downtime.’

The app runs on one instance, so a platform restart or a failure takes it offline until it comes back. Confirm the app is production and matters to the client, add a second instance, turn on Health check, then test that traffic moves when one instance stops responding.

About 2 hours · a second instance doubles that plan’s compute cost

Move App Service plans to Premium v3

‘Paying a premium for yesterday’s technology.’

Premium v3 is available at a lower base rate for comparable capacity, and it makes reservations possible. Check runtimes, slots and outbound dependencies. If the current scale unit can’t run v3, the move means a redeploy, which changes inbound IPs too. Move in a change window, then report cost and health.

About 16 hours · medium risk: inbound and outbound IPs can change

Hours and risk are Spotto’s estimates.

Recurring findings group into five service bundles.

Grouping recurring findings lets you sell them as block time, or deliver them in one go.

Identity and access

Access tied to people and services you can name, manage and revoke when they're no longer needed.

Entra ID sign-in for storage accounts	**Medium job**
Managed identities for apps and services	**Medium job**
A policy to retire shared keys	**Medium job**
Tidy up unused service principals and app registrations	**Size varies**

Network security

Systems reachable from where they should be, and nowhere else.

Private network access for storage accounts	**Medium job**
Private Link for storage accounts	**Medium job**
Network security group rule review	**Medium job**
TLS 1.2 and encryption in transit	**Small job**

Resilience and recovery

Data and systems that can be recovered when something goes wrong.

Zone-redundant storage	**Medium job**
Azure Site Recovery for key workloads	**Medium job**
A second instance for App Service	**Small job**
Backup vault policy coverage	**Medium job**

Database posture

Proof the database is looked after, ready before anyone asks for it.

SQL vulnerability assessment follow-up	**Small job**
Key Vault and application secret hygiene	**Medium job**
Backup, point-in-time restore and retention	**Size varies**
Private endpoints and firewall scope	**Medium job**

Cost and cleanup

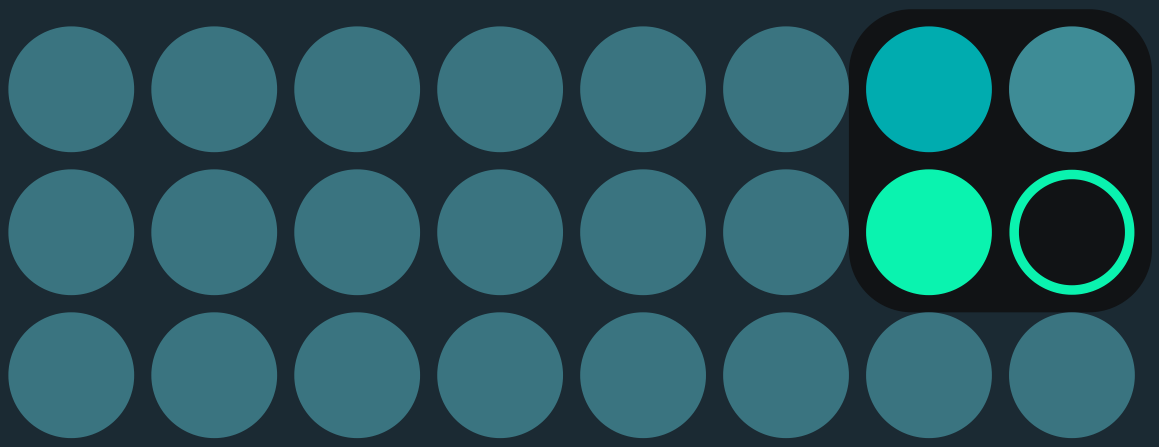
A bill with fewer surprises, and savings that can help fund the rest.

Right-size or shut down underused virtual machines	**Medium job**
Step disks down to the tier the workload uses	**Medium job**
Cost alerts and budgets	**Small job**
Orphaned resource cleanup	**Small job**

3 PUTTING IT TO WORK

From findings to a plan.

Six ways to use what you find, what a well-run environment gives back, and how to run this on your own clients.



SIX WAYS

Using what you find with clients and prospects.

WHAT IT GIVES BACK

The hours it frees up, and what that means for you.

YOUR OWN CLIENTS

How to build the same view on yours.

Six ways to put this to work.

1

Open with cost and quick wins

These are relatively small jobs that can give you a practical starting point with a client. Bundle the quick wins with cost work, then use what you find to shape the next steps.

2

Package by the bundle

Group related findings into one of the five bundles, such as identity and access, or resilience and recovery. Scope related work together, then reuse that approach across the clients who need it.

3

Check before you quote

Some of these settings are deliberate choices. Zone redundancy and Site Recovery add cost, and not every workload needs them. Adding a second App Service instance doubles that plan's compute cost. Agree on the client's requirements before scoping and pricing a change, and record the reasons as context about the client in Spotto, so everyone understands the choices made.

4

Run a themed quarter

Pick one bundle each quarter, such as resilience and recovery, and offer it to every client who needs it. You scope it once, plan delivery as one piece of work, and each client gets a clear focus for the quarter.

5

Bring the benchmark to your client reviews

Highlight to each client which quartile they sit in and what the best-kept environments have already done, then agree a plan to move them towards the best-kept quartile.

6

Bring it to a prospect

Use Spotto in pre-sales to get a full understanding of the environment before the first meeting. You arrive knowing what matters in their Azure environment, so you can have higher-level conversations about outcomes.

A well-run environment gives hours back.

When your team can manage an environment in fewer hours, you have more room to improve margins or deliver more for the client within the team you already have.

WHAT COULD TWO HOURS BACK LOOK LIKE?

Two engineers billing 240 hours a month between them at $200 an hour generate $48,000 in monthly revenue. Using Service Leadership's 48.7% best-in-class services gross margin as a starting point, here's what freeing up two hours a month on each of ten clients could mean.

20 hours

Time back each month

2 hours × 10 clients

$4,000

Potential extra revenue a month

If all 20 hours are sold again at $200 an hour

+8%

Revenue lift

If those hours become additional billable work

48.7% → 52.6%

Illustrative gross margin improvement

If the hours are resold with no extra delivery cost

SHAUN'S TAKE

Amazing if you're selling outcomes and value in your managed service, not time, as that hits the bottom line.

Either way, you now have additional capacity to redirect as you see fit, whether that's additional billable work, service improvement or something else.

What would you do with two hours a month on every client?

WHAT THIS LOOKED LIKE FOR ONE PROVIDER

US$6,000–10,000

of optimisation work every month.

Umbrellar Technology Group, quoted with permission, August 2026.

The hours, revenue and margin figures above are illustrative. The revenue example assumes the 20 hours are sold again at $200 an hour. The margin calculation assumes the additional work is delivered with no extra delivery cost. See Method and scope for sources.

KEY TAKEAWAYS

What this means for your practice.



1 **Your Azure practice is a risk practice.**

Across 417 client subscriptions, security and reliability make up 85.6% of the open work, and cost 3.5%. Cost opens the door. The larger, recurring service is keeping clients' systems secure and running.

2 **Sell the benchmark.**

The best-kept quartile of environments carries 91 or fewer high-impact findings per 100 resources. The least-kept carries 171 or more. Show a client where they sit, and the gap becomes the plan and the proposal.

3 **Turn recurring work into margin.**

The most common findings group into five service bundles. Scope and price a bundle once, then deliver it across your book the same way each time. That repeatability protects your services margin.

4 **Growth comes from reading the environment first.**

Almost every environment carries high-impact work the client hasn't asked about yet. The practices that grow find it first, package it and bring it to the client before anyone asks. That sets the next 12 months of work, for them and for you.

RUN THIS ON YOUR OWN CLIENTS

Everything here describes other providers' clients. Build the same view on yours.

The numbers that will matter most to you are your own. Spend is a reasonable first guess at how much work is open, and a poor guess at how dense it is, so you still need to look inside the environment.

1 **Pick three clients**

Start with three client environments: one large, one mid-sized and one that hasn't had a technical review in a while. That last one often holds the most surprises.

2 **Compare them with the benchmark**

Count high-impact findings per 100 resources and compare with the best-kept quartile, at 91 or fewer. If the mix looks like the one in this report, open with cost and quick wins.

3 **Turn what you find into a roadmap**

Group what you find into the five service bundles, put them in front of the client, and use them to shape the next 12 months of work together.

Connecting an environment uses an Entra app registration, with no agent to install and no user accounts to create. The full list of permissions is in the Spotto docs at **docs.spotto.ai**.

spotto

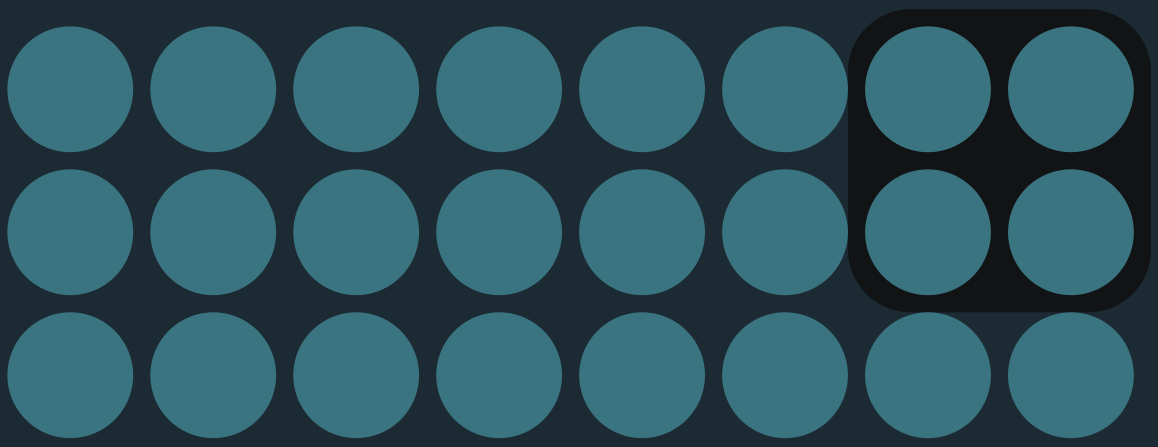
Want to talk about what might be hiding in your clients' Azure?

hello@spotto.ai
www.spotto.ai/hidden-work

A APPENDIX

The detail behind the numbers.

How much clients differ, and where each figure comes from.



HOW MUCH CLIENTS DIFFER

Each client's own mix of work, side by side.

METHOD AND SCOPE

The data, the cuts, and the source of every figure.

Most clients have the same mix of work. A few don't.

What we measured

For each client, we worked out what share of their own open findings falls into each category. Then we lined the clients up side by side.

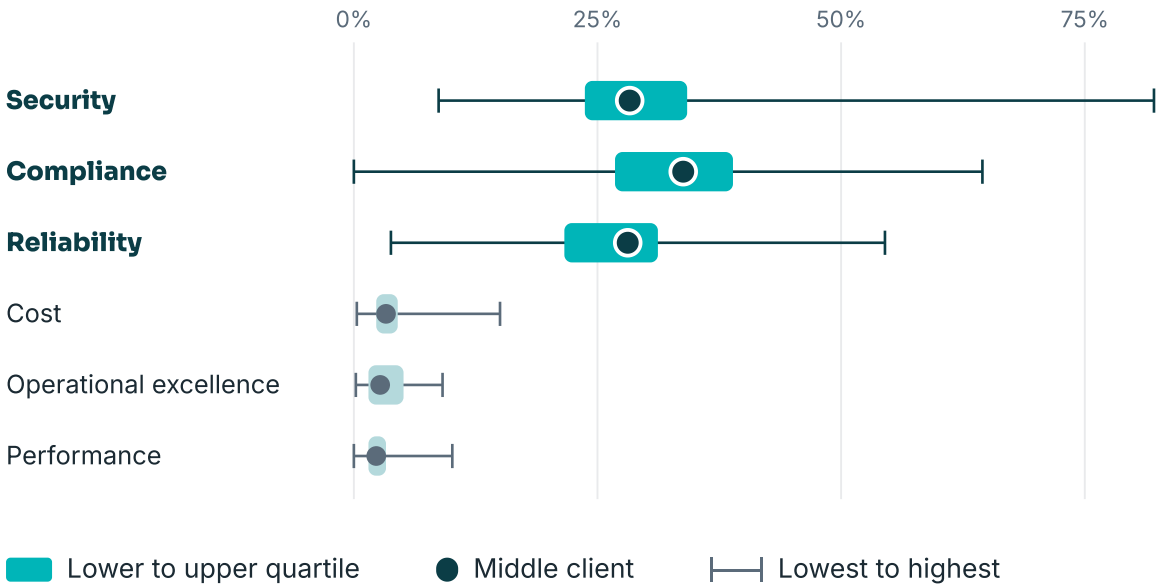
What it shows

A typical client's findings are roughly 30% each in security, compliance and reliability, and a small share in cost. A few clients are very different: one has 82% of its findings in security, another has 15% in cost.

How it helps you

Put one of your own clients next to the typical mix, and look first at wherever it differs most.

SHARE OF EACH CLIENT'S OPEN FINDINGS, BY CATEGORY



How to read it. Line every client up in order. The one in the middle is the median. The two quartile marks fence off the middle half of clients. Lowest and highest are the two extremes, one client each.

Each client counts once, whatever its size. We work out each category on its own, so a client's shares don't add up to 100%.

Where the numbers come from.

DATASET	WHAT IT COVERS	DEPTH	USED FOR
The full pull	417 client subscriptions with open findings, in 66 ANZ client companies	Every finding, with category, impact and effort. Impact and effort leave out policy checks, which carry one blanket rating	Most of the report, including the bundles and the worked examples
The spend read	30-day spend for each subscription, each currency kept separate	Spend next to findings	What a client's Azure spend tells you
The history	Daily counts for 417 subscriptions over up to 12 months; 235 with 90+ days	Open findings by impact and category over time	What good looks like

WHERE EACH FIGURE COMES FROM

Where the open Azure work sits, and the takeaways. Resource instances across 417 client subscriptions with open findings, in 66 ANZ client companies. Each issue is counted once, under the type of work it is (Spotto's classification): security and reliability are 85.6% of open findings and cost 3.5%. If compliance is counted as its own category (adding about 30% to the total), security, compliance and reliability make up 92.5% and cost 2.7%.

What a client's Azure spend tells you. 310 client subscriptions billed in New Zealand dollars with spend above zero, in equal thirds by 30-day spend, and 42 client companies for the per-$1,000 figures. High-impact share leaves out policy checks; security, compliance and reliability are 93–94% in every group. AUD subscriptions show the same pattern.

How big the work is. Grouped recommendation rows, which is a different unit from the resource counts used elsewhere, so the totals don't match. 28,013 of 33,434 rows carry both ratings; policy checks are left out. Effort ratings are Spotto's estimates.

What good looks like. 353 client subscriptions with at least 10 resources, counted as open high-impact findings per 100 resources. The quartiles are cut at 91 and 171 high-impact findings per 100 resources; the typical subscription in the best-kept quartile has 67 and in the least-kept 227 (the quartile medians used in the Highlights). The history covers 235 subscriptions with 90 days or more of daily counts.

WHERE EACH FIGURE COMES FROM, CONTINUED

What good management gives back. An illustration built from provider examples: one expects about 30% of its time back, another spends two hours a month building a report by hand. Best-in-class margin: Service Leadership Index (ConnectWise), December 2025.

How much of the work repeats. High-impact finding rows from the full pull, policy checks left out, counted by resource instance. A kind of finding is one recommendation type; 42 of the 368 are open at 33 or more of the 66 client companies.

How much clients differ. Each of the 66 client companies' own open findings by category, counted by resource, with compliance as its own category (unlike 'Where the open Azure work sits').

Quick wins. 371 of the 417 subscriptions (89%) have at least one high-impact finding rated low effort.

Worked examples. Six jobs from the full pull. Hours and risk are Spotto's estimates.

The five service bundles. Job sizes are Spotto's effort rating for each finding: small is low effort, large is high effort. Where Azure reports the same fix twice, it counts once.

Progress over 90 days. Last 90 days of history; check-change days removed; subscriptions whose resources fell 20% or more left out. 201 subscriptions (192 security, 148 cost). The top quartile by % fall is 'making the most progress', a reference point. Falls can also come from resources or settings changing.

Everywhere. Policy checks (compliance) carry one blanket impact and effort rating, so impact, effort and best-kept figures leave them out.



spotto

spotto.ai

The Hidden Azure Work Index · First edition, 2026