CHARTA LEARNING · LESSON NOTES

Cyber-resilience: Range of Practices

FRM Part II

1 A story: the morning the payments stopped

- A fake invoice email lets an intruder in
- Ransomware locks Halyard's servers
- Defences were sensible, and most attacks were stopped
- New question: can customers still be paid today?

IN THE LESSON

- Picture Halyard Bank on a Monday morning.
- Over the weekend, an employee opened an email that looked like an invoice, and the attachment let an intruder in.
- Early on Monday the intruder launches ransomware, software that encrypts data and demands a payment to unlock it, and starts locking up servers.
- Halyard had firewalls, antivirus software and staff training.
- Most attacks it faced were stopped.
- This one was not.
- So the important question is no longer whether Halyard could have kept the attacker out.
- It is whether the bank can keep paying its customers' salaries today.
- That shift, from keeping attackers out to keeping the business running, is the heart of cyber-resilience.

2 What cyber-resilience means

- Keep carrying out the mission despite cyber threats
- Anticipate threats; withstand and contain incidents
- Recover rapidly, then adapt: learn for next time
- Accepts some attacks will succeed

IN THE LESSON

- Here is the definition.
- Cyber-resilience is the ability of an organisation to keep carrying out its mission by anticipating and adapting to cyber threats, and by withstanding, containing and rapidly recovering from cyber incidents.
- Anticipating means looking ahead at the threats that are coming.
- Withstanding means absorbing an attack without collapsing.
- Containing means stopping it from spreading.
- Recovering means restoring services quickly.
- Adapting means learning, so the next attack meets a stronger bank.
- Notice what the definition does not promise.
- It does not say that no attack will ever succeed.
- It says that critical services continue.

3 Cybersecurity versus cyber-resilience

Security's three properties: confidentiality, integrity, availability

Cybersecurity	Cyber-resilience
Protects data and systems	Keeps critical services going
Secret, correct, available	Withstand, contain, recover
Asks: are we protected?	Asks: do services continue?
Aims to prevent attacks	Accepts some will succeed

IN THE LESSON

- People often treat the two words as the same, so let us set them side by side.
- Cybersecurity is built on three properties: confidentiality, meaning data stays secret from the wrong people, integrity, meaning it stays correct, and availability, meaning it is there when needed.
- Now row by row.
- Cybersecurity protects information and systems, while cyber-resilience keeps critical services going.
- Security keeps data secret, correct and available, while resilience adds the ability to withstand, contain and recover.
- Cybersecurity asks: are we protected?
- Resilience asks: do our critical services continue, even when protection fails?
- And cybersecurity mainly aims to prevent attacks, while resilience accepts that some will succeed.
- Halyard's firewalls were security.
- Its plan for paying salaries on Monday is resilience.

4 Three official initiatives side by side

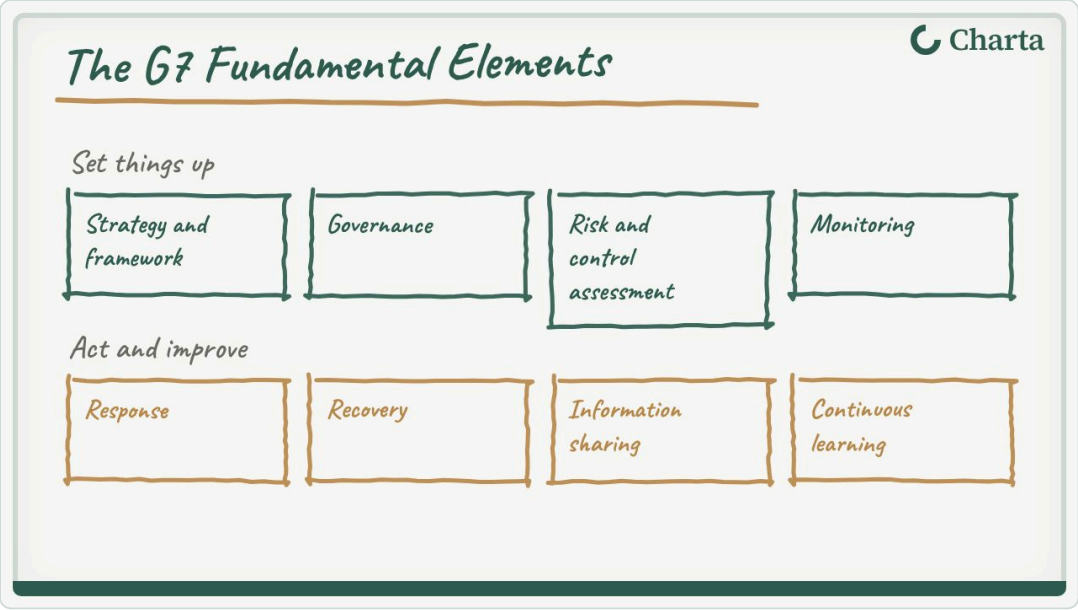
Know who each is aimed at, and its structure

Initiative	Aimed at	Structure
G7 Fundamental Elements	Financial entities	8 elements
CPMI-IOSCO guidance	Market infrastructures	5 categories + 3
FSB Cyber Lexicon	Firms and authorities	Shared vocabulary

IN THE LESSON

- The official sector has built on this idea in several pieces of guidance.
- The G seven Fundamental Elements are aimed at financial entities generally, and they give eight high-level elements.
- The C P M I and I O S C O guidance is aimed at financial market infrastructures, such as payment systems and central counterparties.
- It applies their existing principles to cyber, using five risk management categories plus three overarching components.
- The F S B Cyber Lexicon supplies only shared vocabulary, so that firms and authorities mean the same thing by the same word.
- One point covers all three: none is binding.
- Each is implemented nationally.

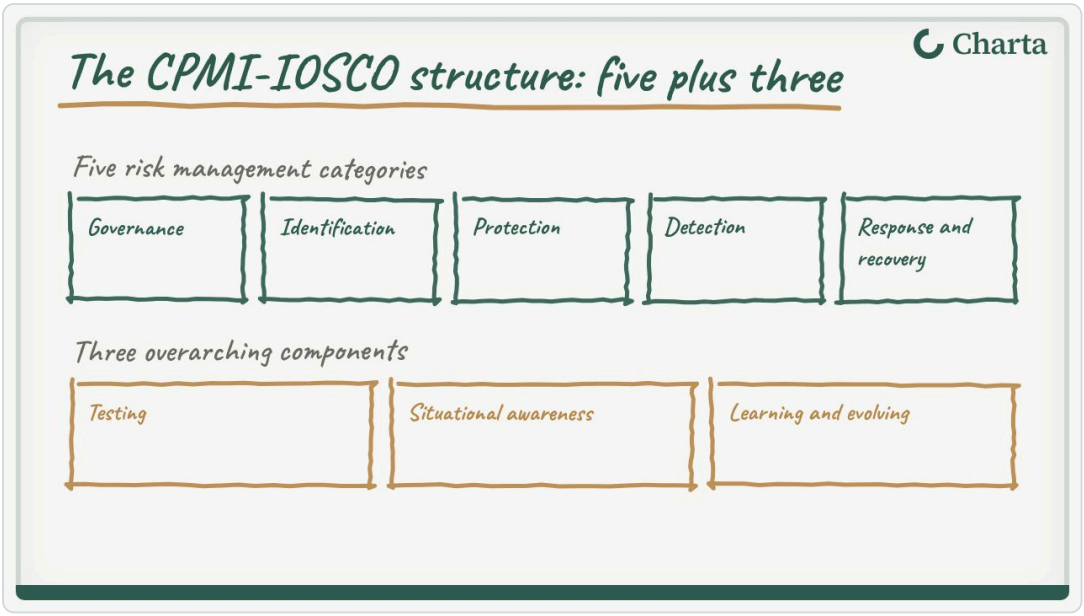
5 The G7 Fundamental Elements



IN THE LESSON

- Here are the eight G seven elements, in two groups of four.
- The first group sets things up.
- A strategy and framework says what the firm is trying to achieve in cyber and how.
- Governance says who is accountable, up to the board.
- Risk and control assessment identifies the important functions and the threats to them, and checks the controls.
- Monitoring watches whether the controls keep working.
- The second group acts and improves.
- Response means acting when an incident happens.
- Recovery means restoring normal operations.
- Information sharing means exchanging threat information with others.
- Continuous learning means feeding experience back into the framework.
- There are no technical rules here: the elements are high-level.

6 The CPMI-IOSCO structure: five plus three



IN THE LESSON

- The C P M I and I O S C O guidance is built differently.
- It has five risk management categories.
- Governance sets the framework and accountability.
- Identification finds the critical functions and the information assets that support them.
- Protection puts controls around them.
- Detection spots incidents quickly.
- Response and recovery restores operations.
- Then come three overarching components that run across all five.
- Testing checks that each part works.
- Situational awareness means understanding the threats the infrastructure faces.
- Learning and evolving means improving the framework as threats change.
- So, five categories, plus three overarching components.

7 Mistakes to avoid: mixing up the structures

- G7: 8 elements, for financial entities generally
- CPMI-IOSCO: 5 categories + 3, for infrastructures
- FSB Lexicon: vocabulary only, not a framework
- None is binding: each country implements it
- Resilience is not just better prevention

IN THE LESSON

- The first trap is mixing up the structures.
- Eight elements belong to the G seven, for financial entities generally.
- Five categories plus three components belong to the C P M I and I O S C O guidance, for market infrastructures.
- The Lexicon is not a framework at all.
- A second trap is to call any of them binding law.
- They are guidance, and each country turns them into its own rules.
- A third trap is to describe resilience as better firewalls.
- Resilience is about keeping critical services going when protection fails.
- When a question names an initiative, ask: who is it aimed at, and how is it built?

8 Who oversees the cyber framework

- Board: approves cyber strategy and risk appetite
- Board needs reporting it can understand
- Not 4,000 alerts, but: 90% on time vs appetite 95%
- Senior management: implements, allocates resources
- Day to day: the three lines of defence

IN THE LESSON

- Now, who is in charge of all this at Halyard?
- Oversight starts with the board of directors.
- The board approves the cyber strategy and the cyber risk appetite, which is how much cyber risk the bank is willing to accept.
- To do that, the board must receive reporting it can understand.
- Four thousand technical alerts tell directors nothing.
- A statement such as, our most critical systems were patched on time ninety per cent of the time, against an appetite of ninety five, tells them whether the bank is inside its appetite.
- Senior management implements the framework and allocates the resources, meaning the people, money and tools.
- Day to day duties then follow the three lines of defence.

9 Three lines of defence for cyber

Operate. Challenge. Assure.

Line	Cyber duties
1st: business and technology	Patch, manage access, monitor
2nd: risk, CISO's function	Policy, challenge, vs appetite
3rd: internal audit	Independent assurance

IN THE LESSON

- Here are the three lines applied to cyber.
- The first line is the business and technology units.
- They own and operate the controls, such as patching, which means installing software fixes, access management, which means deciding who can log in to what, and monitoring of the systems.
- The second line is the risk and information security oversight functions, led by a chief information security officer, or C I S O, or an equivalent role.
- The second line sets policy, challenges the first line, and monitors the risk against the appetite.
- The third line is internal audit.
- It gives independent assurance that the framework and the controls actually work.
- In short: operate, challenge, assure.

10 Example 1: matching duties at Halyard

Which line does each cyber duty belong to?

Installs this month's patches	technology team	**1st line**
Writes admin access policy	CISO's function	**2nd line**
Reports risk vs appetite	CISO's function	**2nd line**
Tests access reviews happen	internal audit	**3rd line**
Approves cyber risk appetite	board	**Board**

IN THE LESSON

- Let us sort five duties at Halyard.
- Installing this month's patches on the servers is done by the technology team, so it is first line.
- Writing the policy on who may have administrator access is done by the C I S O's function, which is second line.
- Reporting to the board that cyber risk is above appetite is also a second line job, because the second line monitors risk against appetite.
- Testing whether access reviews really happen is internal audit, the third line.
- Finally, approving the cyber risk appetite itself sits above the lines, with the board.

11 How authorities differ

Same goals, different rulebooks

Approach	What it means
Cyber-specific rules	A rulebook just for cyber
Extend existing rules	Op risk, IT, outsourcing
Common to both	Board accountable
Common to both	Three lines; show it works

IN THE LESSON

- Supervisors around the world agree on the broad shape, but they write it down differently.
- Some authorities issue cyber-specific rules, a rulebook written just for cyber risk.
- Others do not write a separate rulebook.
- Instead they extend their existing operational risk, information technology and outsourcing requirements so that those rules cover cyber too.
- So one control may sit under different headings in different countries.
- What the approaches have in common is more important for the exam.
- In all of them the board is accountable, the framework is expected to follow the three lines, and the bank must be able to show that its controls work.

12 Mistakes to avoid: oversight under the doer

- CISO reporting to the head of IT: weak challenge
- IT grades its own controls: no independent view
- Audit runs controls it later tests: no assurance
- Fix: 2nd line independent, direct line to board
- Match each duty to the correct line

IN THE LESSON

- The usual trap is placing oversight under the function it must challenge.
- Suppose Halyard's C I S O reports to the head of information technology.
- The C I S O's job is to challenge the technology team, for example about late patches.
- But the person being challenged now controls the C I S O's budget, priorities and pay.
- The challenge is weakened, and bad news may never reach the board.
- The same happens if technology grades its own controls, or audit runs controls it later tests.
- Where the C I S O sits varies between firms, so the fix is not one set reporting line.
- What matters is that whoever performs the second-line challenge is independent of the first line, with direct access to senior management and the board.
- In exam questions, match each duty to the correct line.

13 How supervisors assess cyber-resilience

Five ways a supervisor looks at a bank

Tool	What it involves
Off-site review	Documents and data sent in
On-site inspection	Examiners visit the bank
Thematic review	One topic across many banks
Self-assessment	The bank scores itself
Supervisor-led testing	Authority sets the test

IN THE LESSON

- How does a supervisor judge Halyard?
- It has several tools.
- Off-site reviews examine documents and data that the bank sends in, such as policies, reports and incident logs.
- On-site inspections send examiners into the bank to look at systems and talk to staff.
- Thematic reviews compare the same topic across many banks at once, for example how all large banks handle ransomware.
- Self-assessments ask the bank to score itself against a set of expectations, which is cheap but only as honest as the bank.
- And sometimes there is supervisor-led testing, where the authority sets or oversees a test of the bank's defences.
- Most supervisors combine several of these.

14 The ladder of security tests

Scan, pen test, red team climb; tabletop tests decisions

Test	What it does
Vulnerability scan	Finds known weaknesses
Penetration test	Tries to exploit them
Tabletop exercise	Rehearses decisions, no systems
Threat-led red teaming	Imitates real attackers

IN THE LESSON

- Institutions test themselves with a ladder of methods.
- The scan, the penetration test and red teaming climb, each going further than the one below, while the tabletop exercise tests decisions.
- A vulnerability scan is an automated sweep that finds known weaknesses.
- A penetration test goes further.
- Testers try to exploit the weaknesses, to see whether they can actually get in.
- A tabletop exercise is different in kind.
- Managers sit around a table and rehearse their decisions and communication during an imagined incident, without touching any system.
- At the top, threat-led red teaming imitates real adversaries.
- It uses threat intelligence about how actual attackers behave, aims at the critical functions, and is often run without telling the defenders, so that detection and response are genuinely observed.

15 Example 2: fit the test to the objective

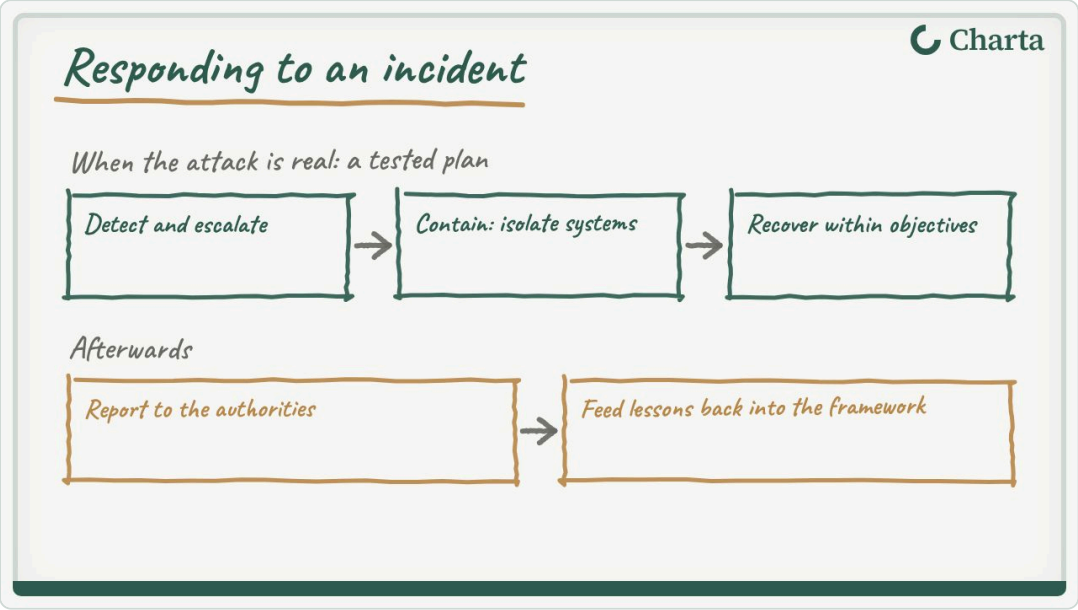
Halyard's four objectives, and the right test for each

Objective	Test
Which servers miss patches?	Vulnerability scan
Can an outsider break in?	Penetration test
Will executives decide well?	Tabletop exercise
Would we spot a real attack?	Red team, unannounced

IN THE LESSON

- Now the exam skill: choose the test that fits the stated objective.
- Halyard has four questions.
- First, which of our servers are missing known patches?
- That calls for a vulnerability scan.
- Second, can an outsider actually break in through our online banking portal?
- That calls for a penetration test.
- Third, would our executives decide quickly, and tell customers and the supervisor the right things, during a ransomware crisis?
- That is a tabletop exercise, because the objective is decisions and communication.
- Fourth, would our security team notice a realistic attack on payments?
- That calls for threat-led red teaming, unannounced, because only then is detection genuinely observed.

16 Responding to an incident



IN THE LESSON

- Testing prepares the bank.
- Incident response is what happens when the attack is real, and it needs tested plans, not plans written on the day.
- Look at the boxes.
- First, detect the incident and escalate it along a set path, so the right people know within minutes.
- Second, contain it, by isolating infected systems so the damage stops spreading.
- Third, recover, restoring services within the recovery objectives agreed in advance.
- Then the bank reports the incident to the authorities, as its supervisors require.
- Finally, the lessons are fed back into the framework, so controls, plans and tests improve.
- That is the adapting part of resilience.

17 Example 3: Halyard's incident timeline

Hours from attack: detected 6, contained 9, payments back 13; payments ran until 6

Time to detect	attack start to hour 6	**6 hours**
Time to contain	9 − 6	**3 hours**
Time to recover	13 − 9	**4 hours**
Payments down	13 − 6	**7 hours**
Recovery objective: 8 hours	7 is within 8	**Met**

IN THE LESSON

- Let us put times on Halyard's ransomware attack, in hours from when it began on Monday.
- The security team detects it at hour six, so the time to detect is six hours.
- It isolates the infected servers and contains the attack at hour nine, so the time to contain, counted from detection, is three hours.
- Payments are restored at hour thirteen, so the time to recover, counted from containment, is four hours.
- Here each time is counted from the previous step; some banks count every time from detection, so always check the starting point.
- Payments kept running until hour six, when Halyard shut them down to stop the spread, so they were down from hour six to hour thirteen, which is seven hours.
- A recovery objective is measured from when the service stops, not from when the attack starts.
- The recovery objective for payments was eight hours.
- Seven is within eight, so the objective was met, and Halyard reports the incident and its timeline to the supervisor.

18 Three groups of metrics

Halyard's dashboard, and a phishing test of 2,000 staff

Group	Examples
Protection	Patch rate; phishing clicks
e.g. phishing test	160 ÷ 2,000 = 8% clicked
Detection and response	Time to detect, contain, recover
Incident counts	Number by type
Note	No standard set exists

IN THE LESSON

- Halyard's metrics fall into three groups.
- Protection metrics show how well the doors are locked.
- Examples are patch compliance, the share of systems patched on time, and phishing-test results.
- In a phishing test the bank sends staff fake scam emails.
- If one hundred and sixty of two thousand staff click, the click rate is eight per cent.
- Detection and response metrics measure the time to detect, the time to contain and the time to recover, as in the timeline we just saw.
- Incident counts form the third group: how many incidents of each type occurred.
- No standard set of metrics exists.
- Each bank chooses its own, so comparisons across banks need care.

19 Example 4: an average spoiled by one incident

Hours to detect five incidents: 2, 3, 4, 5 and 96

Total hours	2 + 3 + 4 + 5 + 96	**110**
Mean	110 ÷ 5	**22 hours**
Median: the middle value	2, 3, 4, 5, 96	**4 hours**
Report the outlier itself	one incident	**96 hours**

IN THE LESSON

- Metrics can mislead in two ways.
- The first is an average distorted by an outlier, a value far from the rest.
- Halyard had five incidents last quarter.
- They took two, three, four and five hours to detect, and one took ninety six hours.
- Add them up and you get one hundred and ten hours, so the mean is twenty two hours.
- But four of the five incidents were found in five hours or less.
- The median, the middle value when they are sorted, is four hours, which describes the typical incident far better.
- Still, report the ninety six hour case on its own, because the board needs to hear about it.

20 Example 5: the wrong denominator

1,800 servers patched. Register says 2,000; a full scan finds 2,400

Reported rate	1,800 ÷ 2,000	**90%**
Servers missing from register	2,400 − 2,000	**400**
True rate	1,800 ÷ 2,400	**75%**
Gap hidden by the denominator	90% − 75%	**15 points**

IN THE LESSON

- The second trap is a compliance rate with the wrong denominator, the number at the bottom of the fraction.
- Halyard's technology team reports that one thousand eight hundred servers are patched, out of the two thousand servers in its register.
- That is ninety per cent, close to the appetite.
- But a network scan finds two thousand four hundred servers in total.
- Four hundred servers were never in the register, so they were never patched and never counted.
- The true rate is one thousand eight hundred over two thousand four hundred, which is seventy five per cent.
- So the wrong denominator hid a gap of fifteen percentage points.
- Forgotten servers are often the least protected.
- So always ask what the denominator is, and whether the asset inventory is complete.

21 Mistakes to avoid: tests and metrics

- Tabletop: decisions only; scan: no exploitation
- A red team the defenders expect observes little
- Averages hide outliers; check every denominator
- Fewer incidents counted may mean weaker detection
- Choose the test or metric that fits the objective

IN THE LESSON

- A tabletop exercise does not test technology, because it touches no systems.
- A vulnerability scan lists known weaknesses but does not show whether they can be exploited.
- A red team that everyone knows about does not genuinely observe detection, because the defenders are ready and waiting.
- For metrics, beware averages distorted by outliers, and rates with the wrong denominator.
- Be careful, too, with incident counts.
- A falling count may mean fewer attacks, or it may mean the bank is detecting fewer of them.
- The rule for every question is simple: choose the test or metric that fits the stated objective.

22 A story: one warning, many banks

- Halyard finds the attackers' address and file
- It shares them the same afternoon
- Brookmere blocks the address and stays clean
- Defend collectively: block it before it spreads

IN THE LESSON

- Back to Halyard's bad week.
- While cleaning up, its analysts find the address of the server the attackers used, and a digital fingerprint of the malicious file.
- Halyard shares these with other banks the same afternoon.
- A small savings bank, Brookmere, sees the same fake invoice arriving in its staff inboxes the next morning, blocks the address, and is never infected.
- That is the case for information sharing.
- It lets institutions defend collectively, because a threat seen at one firm can be blocked at others before it spreads.
- The question is how sharing actually happens, and what gets in its way.

23 Where cyber information is shared

Four channels for sharing

Channel	Example
Sector bodies	Information sharing centres
National response teams	Government-backed teams
Public-private channels	Firms and authorities
Bilateral exchanges	Trusted peers call directly

IN THE LESSON

- In practice, sharing happens through four kinds of channel.
- First, sector bodies such as information sharing and analysis centres, where member firms in one industry, for example finance, pool what they see.
- Second, national computer emergency response teams, government-backed teams that coordinate responses to incidents across a country.
- Third, public-private channels with authorities, where supervisors, central banks or security agencies exchange information with firms.
- Fourth, bilateral exchanges between trusted peers, for example the security heads of two banks who know each other and call directly.
- Halyard used its sector body, and phoned a trusted peer.

24 What is shared, and who gives and gets

From precise indicators to big-picture trends

Item	Detail
Technical indicators	Bad addresses, file hashes
Strategic analysis	Trends in attacks
Large institutions	Own intelligence; give a lot
Smaller institutions	Rely on sector bodies

IN THE LESSON

- What is shared ranges widely.
- At one end are technical indicators, the precise details of an attack, such as malicious internet addresses and file hashes.
- A file hash is a short code computed from a file, like a fingerprint, that lets another firm recognise the same malicious file.
- At the other end is strategic analysis of trends, such as which kinds of attackers are targeting banks this year and why.
- Who contributes differs by size.
- Large institutions often run their own threat intelligence teams and contribute heavily.
- Smaller institutions, like Brookmere, rely on sector bodies to gain intelligence they could not afford to produce themselves.

25 Barriers to sharing, and the answers

Pair each barrier with its practical mitigation

Barrier	Mitigation
Liability, rivalry, borders	Legal clarity from authorities
Confidentiality, data rules	Sanitise: strip personal data
Reputational fear	Anonymise the source
Weak trust in onward use	Marking rules on redistribution
Free riding	Reciprocity: give to get

IN THE LESSON

- Sharing faces real barriers, and each has a practical answer.
- Legal liability, competition concerns between rivals, and cross-border differences in law make firms fear that sharing is illegal somewhere.
- The answer is legal clarity from the authorities about what may be shared.
- Confidentiality and data protection rules forbid passing on customers' personal data, so firms share sanitised information, with personal details removed.
- Reputational fear is met by anonymised sharing, which hides the source.
- Weak trust that a report will stay within the group is met by marking rules that limit redistribution.
- And free riding is met by reciprocity.
- Pair each barrier with its practical mitigation.

26 Judging how well sharing works

- Works only with reciprocity: members must give
- Large banks give most, and gain a safer system
- Small banks gain most: intelligence they can't build
- Speed matters: late warnings are history
- Weak points: trust, legal barriers, free riders

IN THE LESSON

- Sharing works only with reciprocity.
- If every bank waits for others to share first, nothing is shared, so a sector body has to make giving part of the deal of being a member.
- The arrangement suits different institutions differently.
- Large banks give more than they get, but they gain from a safer system, since an attack on a weak peer can still hurt them through payments and confidence.
- Small banks gain the most, because they receive intelligence they could not produce.
- Speed matters too: a warning that arrives after the attackers have moved on is history, not defence.
- The weak points are trust, the legal barriers, and members who only take.

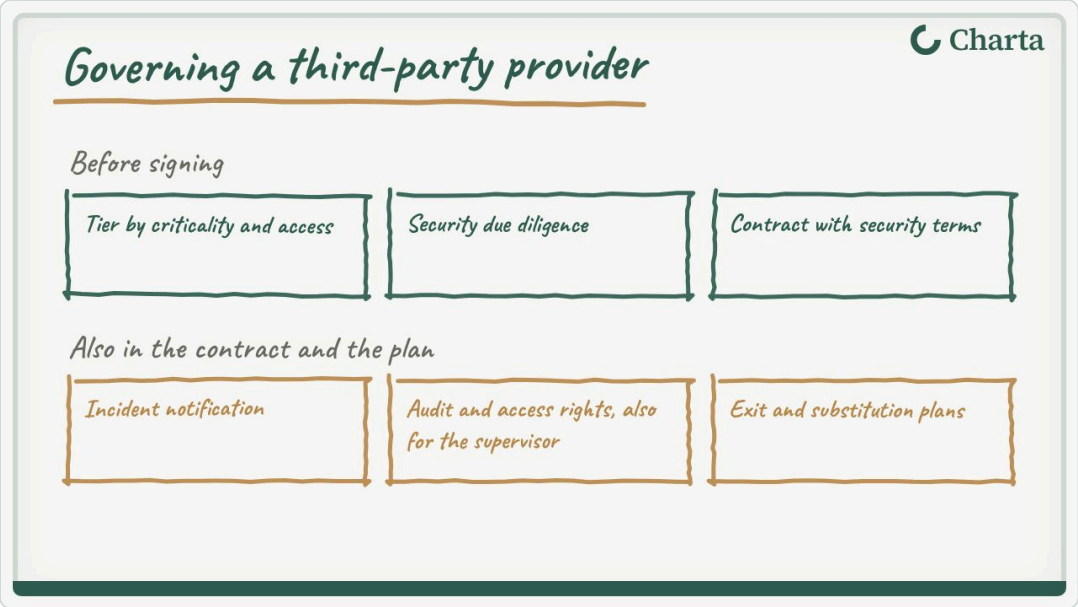
27 A story: the risk that lives elsewhere

- Relies on cloud, software, data-centre, payment firms
- So resilience must reach beyond the bank
- Outsourcing does not transfer accountability
- Halyard stays responsible to its supervisor

IN THE LESSON

- Halyard does not run everything itself.
- Its payments run on a cloud provider's computers.
- Its card transactions go through a payment processor.
- Its core banking software comes from a vendor, and its backup data sits in a third party's data centre.
- Banks depend on cloud, software, data-centre and payment providers, so cyber-resilience must extend beyond the institution's own walls.
- The governing principle is this.
- Outsourcing does not transfer accountability.
- If the cloud provider is breached and Halyard's customers cannot be paid, Halyard remains responsible for the risk and for meeting its supervisors' expectations.
- The supervisor will not accept, it was the provider's fault, as an answer.

28 Governing a third-party provider



IN THE LESSON

- Because it stays accountable, Halyard governs each provider through a set of practices.
- Look at the boxes.
- First, tier the providers by risk, according to how critical the service is and how much access they have.
- Second, carry out security due diligence before signing, checking the provider's controls.
- Third, write contracts that contain security requirements.
- The second row shows what the contract and the plan must also cover.
- Incident notification obliges the provider to tell Halyard quickly when something goes wrong.
- Audit and access rights let Halyard, and its supervisor, examine the provider.
- And exit and substitution plans say how Halyard would move the service to another provider, or back in-house, if the provider failed or had to be dropped.

29 Example 6: tiering Halyard's providers

Annual contract value, and what each provider touches

Cloud for payments: $3m	critical, holds the data	**Tier 1**
Facilities firm: $8m	supervised building access only	**Tier 3**
Patch tool vendor: $0.2m	admin access to servers	**Tier 1**

Biggest contract, lowest tier: size does not set the tier

IN THE LESSON

- Halyard tiers three providers, with made up annual contract values.
- The cloud provider that runs payments costs three million dollars a year.
- Payments are critical and the provider holds the data, so it is tier one, the closest oversight.
- The facilities firm that cleans and maintains the offices costs eight million dollars a year, the biggest contract.
- But it runs no critical service, it has no access to systems or data, and its building access is supervised, so it is tier three.
- A small vendor whose patching tool has administrator access to every server costs only two hundred thousand dollars a year.
- So it is tier one.
- Criticality and access, not contract size, drive the tier.

30 Fourth parties and concentration

- Fourth parties: the provider's own subcontractors
- Governed by disclosure and flow-down of terms
- Concentration: many depend on one provider
- Or on one hidden upstream dependency
- Then one failure spreads widely

IN THE LESSON

- Two risks hide behind the providers.
- The first is fourth parties, the subcontractors on which a provider itself relies.
- Halyard's cloud provider might rent space in someone else's data centre.
- Halyard has no contract with that fourth party, so governance works through disclosure, the provider telling Halyard which subcontractors it uses, and through flow-down, the provider passing Halyard's security requirements down into its own contracts.
- The second is concentration risk.
- It arises when many institutions, or many of one institution's services, depend on the same provider, or on the same hidden upstream dependency.
- Then one failure spreads widely, hitting many services or many banks at once.

31 Example 7: a hidden concentration

Halyard's 10 critical services, and who they run on

On Provider A	from the register	**6 of 10**
On Provider B	from the register	**3 of 10**
In-house	Halyard's own	**1 of 10**
Both A and B use data centre D	6 + 3 = 9; 9 ÷ 10	**90%**

IN THE LESSON

- Let us find a hidden concentration at Halyard.
- The bank has ten critical services.
- It put six on cloud Provider A and three on cloud Provider B, and runs one in-house.
- On paper no provider has more than six of the ten.
- But when Halyard asks both providers to disclose their subcontractors, it learns that both rent their computers in data centres run by the same operator, call it D.
- So nine of the ten critical services, ninety per cent, depend on D.
- If D fails, nine services fail together.
- The fix is to bring this fourth-party dependency into the exit and substitution plans, for example moving some services to a provider that does not use D.

32 Supervisors and the traps

- Supervisors mainly hold the bank to account
- Some can examine critical providers directly
- Traps: thinking accountability moves; tier by size
- Ignoring fourth parties behind the provider
- Two providers may share one hidden dependency

IN THE LESSON

- Who supervises all this?
- Authorities mainly hold the regulated institutions to account.
- Halyard's supervisor examines how Halyard governs its providers, which is why the audit and access rights must also cover the supervisor.
- Some authorities can also examine critical service providers directly, because a failure at one large provider could hit many banks at once.
- Now the traps.
- First, thinking that outsourcing transfers accountability.
- It does not.
- Second, tiering providers by contract size.
- Criticality and access drive the tier.
- Third, stopping at the third party and ignoring the fourth parties behind it.
- Fourth, assuming that using two providers removes concentration, when both may share a hidden upstream dependency.

33 Recap: the main ideas

- Resilience: keep the mission going despite attacks
- G7 8; CPMI-IOSCO 5 + 3; FSB vocabulary; non-binding
- Board approves appetite; operate, challenge, assure
- Scans to red teaming; beware outliers, denominators
- Share with reciprocity; accountability stays home

IN THE LESSON

- Let us pull it together.
- Cyber-resilience is keeping the mission going: anticipating and adapting to threats, and withstanding, containing and recovering from incidents.
- It is wider than cybersecurity.
- The G seven gives eight elements for financial entities, C P M I and I O S C O give five categories plus three components for market infrastructures, and the F S B Lexicon gives vocabulary, all non-binding.
- The board approves strategy and appetite, and the first line operates, the second challenges and audit assures.
- Tests climb from scans to threat-led red teaming, and metrics must survive outliers and bad denominators.
- Sharing needs reciprocity and practical answers to its barriers.
- And outsourcing never transfers accountability, so providers are tiered by criticality and access.

34 Try three quick questions

1. Times to contain: 1, 2, 2, 3 and 42 hours. Mean, median, which is typical?
   Answer: Mean 50 ÷ 5 = 10 hours; median 2 hours, the typical one.
2. 950 of 1,000 registered servers patched; a full scan finds 1,250. True patch rate?
   Answer: 950 ÷ 1,250 = 76%, not 95%.
3. $12m cleaning contract, or $0.1m vendor with admin access to core systems: which is tier 1?
   Answer: The $0.1m vendor: criticality and access set the tier.

IN THE LESSON

- Now try three quick questions.
- Pause, think, then listen.
- Question one.
- Five incidents took one, two, two, three and forty two hours to contain.
- What are the mean and the median, and which describes the typical incident?
- The total is fifty, so the mean is ten hours, the median is two hours, and the median describes the typical incident.
- Question two.
- A team reports nine hundred and fifty of the one thousand servers in its register patched, but a full scan finds one thousand two hundred and fifty servers.
- What is the true patch rate?
- Nine hundred and fifty over one thousand two hundred and fifty is seventy six per cent, not ninety five.
- Question three.
- A twelve million dollar cleaning contract or a one hundred thousand dollar vendor with administrator access to core systems: which is tier one?
- The small vendor, because access and criticality set the tier.