DORA, AI and Cyber Resilience: Full Webinar Transcript

Section 1 — Introduction

Setting the stage: why this conversation cannot wait

Bruno: Hello everybody, and welcome to this webinar. I am Bruno Guinchard, and I will be your moderator today.

Before we begin, a few words about me. I spent the last 18 years at UBS, a Swiss bank, where I held several roles in IT, including quality assurance, testing, compliance, and more recently, AI governance.

Let me start by setting the context for today's discussion and making sure we are all aligned on what has happened in this space over the last few months.

The story really began in April, when Anthropic announced a new version of Mythos. Within a couple of weeks, the model was reportedly able to identify more than 10,000 vulnerabilities across major operating models and browsers, which created a major impact.

In response, Anthropic decided not to launch it as a commercial release, but instead to create Project Glasswing, through which a limited number of companies were invited to test the model and assess the strength of their own defenses. Initially, only U.S. companies were included, such as major hyperscalers like Google, Amazon, and Microsoft.

At the same time, OpenAI released ChatGPT 5.5 Cyber and also gave a limited number of global companies, including some in Europe, the opportunity to test the model and evaluate their defenses.

More recently, just last week, Anthropic announced that it was extending Project Glasswing to more companies, especially in Europe, although the names of the companies have not yet been disclosed. There is also some ongoing discussion with the European Union about access conditions.

So, the pace of change is quite fast. The question for today is: given everything that is happening, what is the main impact for financial institutions and for leaders? What should we do?

We will discuss this through four lenses. First, trust: how is trust in our clients and institutions affected by what is happening? Second, continuity: can we maintain business continuity in the face of new threats? Third, accountability: what role will regulators play? And finally, what next: what actions do we need to take now to make sure we are prepared for what is coming?

To explore these questions, I am joined by a great panel. We have Andrei Kucharavy, founder of Surelio and professor in machine learning and software security. We also have Daniela Sozzi, founder of DNYC, who works closely with the fintech ecosystem and brings deep experience in that area. And we have Caroline Perriard, co-founder of Stratedge, who also has significant consulting experience. With that, let's start the discussion.

Section 2 — Trust

New AI models, new risks: can financial institutions still be trusted?

Bruno: First, I'd like to start with you, Andrei, and ask what is really changing here. We see models like OpenAI's and Mythos entering the field. Are these models a true game-changer, or is this simply a trend that has been building for a while? From a technical perspective, how do you see it?

Andrei: Thank you, Bruno, and thank you for the question.

To give some background to the audience, I have been researching the use of generative AI in cybersecurity with the Swiss Cyber Defense Campus since 2020, so for almost six years now.

One thing we already observed back then was that even models that are now considered old or outdated, such as GPT-3, were very effective at finding vulnerabilities when used by competent people. For example, a researcher at Pennsylvania used GPT-3 in 2021 to find over 200 vulnerabilities in open-source projects, whereas the best specialized tools at the time could identify only about 100.

So we already knew that large language models had strong capabilities for vulnerability discovery. However, there were two important barriers until recently.

First, the person operating the tools needed to be highly competent, know where to point them, and understand how to turn the vulnerability into an actual attack.

Second, there was the problem of weaponization. Finding a vulnerability is not enough on its own. You still need to figure out how to exploit it, whether it can be combined with other weaknesses, and how it can be used effectively.

That is where coding-capable LLMs started making a difference, especially from 2022 onward, and more significantly in 2023.

What we started to see at the end of 2024 and into 2025 was that fairly advanced models, not only from Anthropic and OpenAI but also code-assistance models from China and Mistral, were able to help put all these ideas together. They still required some human guidance, but they could accelerate the process substantially.

Starting from the end of 2025, models like Claude Opus, and to some degree GPT-5, began automating much of that reasoning process. By autumn 2025, in CTF competitions — capture-the-flag challenges used by hackers in controlled environments — teams were relying heavily on Claude Opus to solve challenges that would normally take a full day in just a few hours.

The real change that Mythos brought was that it was optimized for attack. It included offensive capabilities by design. There is a lot of debate in both the cybersecurity and AI safety communities about whether that was a good idea, because while we now have very strong automated attack capabilities, that is not only due to Mythos. GPT-5.5 has similar capabilities, and Microsoft has also released models with the same kind of power.

The problem is that we do not have the defensive tools ready to match them. So we now have a serious imbalance between attackers and defenders. Mythos and the hype around it brought that to public attention, but the gap was already there from late 2025 and is likely to get worse over the next 18 to 20 months before it improves.

Bruno: Thank you for that excellent explanation. So you are saying there is a real new trend emerging, and financial institutions have always been a front-line target in this space.

Daniela, where do you think financial services are more exposed than other industries in this context?

Daniela: As you said, financial services have always been targeted by criminals because they store money or facilitate the movement of money.

What we are seeing now, however, is that these new tools, combined with the fact that financial institutions are the connective tissue of the economic environment, create a much broader threat. Malicious actors, sometimes state-sponsored, may use attacks on financial institutions not just to steal money, but to disrupt entire economies. And they have the potential to do exactly that.

I think we also need to consider the distinction we often make between Tier 1 banks and Tier 2 financial institutions. Anthropic seems to have prioritized Tier 1 banks because they are larger and therefore believed to need priority in cybersecurity testing. But if we look at small and medium-sized enterprises —

the thousands of businesses that hold economies together and are often the driving force in many countries — they rely heavily on Tier 2 financial institutions, such as Sparkassen, local savings banks, cooperative banks, and similar providers.

These are smaller institutions, there are many more of them, and they are often perceived as having fewer defense mechanisms than Tier 1 banks. This is really the concern everyone is focused on: central banks, regulators, and institutions alike.

The reason is simple. What they want to preserve is trust in the system. If a malicious actor succeeds in an attack, the real damage is not only to the individual institution. The greater damage is to trust in the system itself. And that is what we must prevent.

Bruno: Yes, as you said, trust is key here. This is not just an IT issue; it is a global issue for the entire organization.

I assume you also work with a lot of boards. What are the main topics you are discussing with them right now, and what fears are they expressing on this subject?

Caroline: Thank you, Bruno. I think a lot of topics are being raised, but I would focus on three in particular.

The first is compliance: how do we ensure that AI models and new technologies comply with regulation? Regulation is evolving, and different jurisdictions around the world rely on different frameworks. Some are sector-specific, some are very stringent, and some rely on existing legislation. In Europe, Switzerland, and other countries, regulation is still being developed, with many guidelines aimed at understanding how the technology works and what the risks are.

The second is auditing and documentation. How do we audit and document AI decision-making for regulators? Financial institutions need to document their processes properly and ensure they have the right audit trail.

For example, if an AI-driven solution is used to decide whether someone receives a credit limit, and that solution contains bias, the result may be discriminatory. That could violate anti-discrimination law, and the company would need to show that it had taken appropriate measures to prevent that bias.

The issue is how to audit and document before harm occurs. If the data feeding the model reflects historical inequalities — for example, data showing that women are paid less than men — then the tool may infer that a woman has fewer means than a man and assign a lower credit limit. In the end, that becomes discriminatory.

The third issue is trust in the solution itself, especially when it is used for critical financial decisions. How do we mitigate systemic risk?

Mitigation requires a lot of technological and IT knowledge. At the end of the day, companies have to ask whether the tool is worth the extra operational

burden, because it may require extensive pre-testing, verification, and other safeguards to be used responsibly.

So I would say the main fears are reputation risk and loss of trust. In the financial sector, there is also the fear of losing licenses or even the ability to operate. That is a very serious burden for any institution.

Bruno: Thank you. Personal exposure and compliance risk are clearly key issues for all institutions.

Daniela, coming back to your fintech experience: we know fintechs are often more agile and can move faster than Tier 1 banks, but they may also have less infrastructure and less capacity to protect themselves. From your point of view, does the fintech ecosystem see what is happening as a threat, an opportunity, or both?

Daniela: It is a bit of both, and I will focus particularly on European fintech, which is where I am closest to the market.

From a cybersecurity perspective, there is an opportunity, because the largest fintechs in Europe are often perceived as more agile and potentially more resilient to cyberattacks. I stress that this is still a perception. The real test would come if an incident actually occurred. But at the moment, this is definitely seen as an opportunity to win market share from incumbents and more traditional players.

At the same time, it is also a threat, because fintechs do not operate in a silo or in a parallel world. They are interconnected. In many cases, they resell services or asset management products produced by traditional financial players. Because they are part of the same ecosystem, an attack on the ecosystem is also a threat to fintechs. And they know that. That is why this is both a threat and an opportunity, and why it is fully reflected in their risk assessments.

Section 3 — Continuity

When the attacker has AI too: can you still keep the lights on?

Bruno: Okay, interesting. So we have discussed the first part, focused on trust and the key relationship between financial institutions and clients, as well as compliance risk and reputational risk.

Let us move now to the question of continuity. For me, continuity is about how we make sure that, in a crisis or difficult situation, we can still run the business.

With everything that is happening now, we need to think about the impact on our operational processes and on the people involved. So let us discuss that, starting with the damage-control approach: the risk management approach we currently use, and how we control threats at this stage.

Do you see any change in the approach compared with what is already in place today? Do we need to change the way we work to stay prepared? Andrei, what is your view?

Andrei: Yes, I think it is a bit of both.

On the one hand, traditional cybersecurity approaches do not suddenly stop working in the face of AI-powered attackers. There have been many headlines about vulnerabilities being discovered by Mythos, but there was also an attempt to use Mythos to scan a common tool called CURL for vulnerabilities. Mythos found zero vulnerabilities. It actually found one issue, but it was more of a typo than a real vulnerability.

The reason was that the developer of the tool had already put classical cybersecurity measures in place. They had extensive testing, security audits, and other reviews, including testing with older models. If you are already following those best practices thoroughly, there is a good chance you are still very difficult to crack, even for these new offensively tuned tools.

So if you have strong classical cybersecurity practices, you are already in a good position.

However, no defense is perfect. There is no such thing as an impenetrable fortress. Eventually, as the tools get better and attackers become more persistent, something will get through. Someone will get onto your network, acquire administrative privileges, and you need to be ready for that.

One of the major changes is speed. Once an attacker is inside, lateral movement is now happening so quickly that it is becoming impossible to control with classical tools alone. There was a recent compromise of a tool called LightLLM, where an LLM-based security scanner running on a repository was compromised, secrets were stolen, and within only about 30 minutes a version of the tool with a payload was posted and pushed into production. Seventeen minutes later, it was updated again.

Historically, we would have expected that transition to take at least a couple of days, and usually a couple of weeks, even for advanced attackers. Now we are talking about minutes.

The same thing is happening when a patch is released. We can now use LLMs to reverse-engineer the majority of the attacks that the patch addresses and create a proof of concept exploit for about 75 percent of them in under an hour, and 90 percent within 24 hours.

That means that if your previous practice was to wait a little before patching, to make sure nothing breaks in the company, that is no longer viable. And if your incident response depends on humans moving manually, it is too slow.

You now need something that detects the intrusion and starts stopping or isolating parts of the network within minutes of the initial breach. We are moving into a velocity that is too fast for humans, and that will probably require agents.

But then we run into another issue: if you deploy agents for defense, those same agents can also destroy data, damage the network, or delete what has already happened. So we need to start thinking now about governance and about limiting the capabilities of those defensive agents in the network, even if the best tools are only just beginning to emerge.

In my opinion, some of the really good tools will arrive in the next 12 to 15 months, and others may be available in just a few months.

Bruno: Thank you. The key point is velocity. The approach may still be the same in principle, but we need to be much quicker. That will definitely have an impact on operations.

So if we look at financial institutions, what are the main operational impacts of these new models? Daniela, what is your view?

Daniela: Building on what Andrei just described, the first issue is really processes.

We are seeing a shift from being reactive — waiting for an incident to happen and then asking what to do next — to focusing much more on prevention. That includes the technical tools Andrei just described, but also the separation of decision-making from automated solutions so that attacks can be prevented and responded to as fast as they unfold.

This changes operations as well. I would borrow a concept from military strategy here. What do you do if your opponent no longer follows the code of war or the usual rules of engagement? You have to adapt your strategy.

If it becomes a guerrilla-style environment, you cannot simply deploy your battalion and rely on numbers anymore. You need to be more agile and more responsive, and operations need to be designed accordingly.

That requires a lot of training. It is not something that can be improvised. It is not something you can wait to learn after an incident happens. It has to be prepared in advance.

Historically, operations in financial services have often been more hands-oriented than brains-oriented, with people waiting for instructions from the top or following checklists. That will not work in today's environment. There has to be a shift.

This also affects headcount. It is no longer just about numbers. It is about the quality of the people available when an incident occurs, and all hands need to be ready.

Those people need a skill set that is different from the past — one that fits the present and the future. Empowerment follows from that. There is no time to wait for instructions, and no time to wait for a traditional analysis process.

So processes, operations, and headcount all need to be adapted. Some people may assume that this means fewer people, but that is not the point. It is not about reducing headcount. It is about having the right skill set available at the right time.

Bruno: So we need more qualified people, and we need processes that can adapt to the new pace we are seeing. That means a lot of operational change.

On the one hand, we need more action to protect against cyberattacks. On the other, we still need to keep the business running. So we need to find the right trade-off between business continuity and cyber defense. Caroline, how do we define that trade-off, and how do we make the decision?

Caroline: Yes, thank you. I agree with Andrei, and I also concur with Daniela. She raised many important initiatives that should be taken.

The balance is definitely difficult. I think we need to consider the budgets companies have, as well as the critical data and systems they hold. There is no one-size-fits-all solution for deciding how much to reinforce cyber defense.

It needs to be assessed in relation to the consequences of stopping business, delaying services, or losing access to essential data. We need to ask whether the critical functions are being attacked or whether it is something on the side.

In the past, I have been involved in many data leakage incidents, and the main challenge was always: where is the data, and what data has been accessed? I now wonder whether AI tools will help us identify that faster and more easily. I am not sure.

If your house is not in order — if the foundation and the cellar are a mess — then it will also be difficult to layer on top all the capabilities that these tools promise. We cannot simply assume that we are protected while keeping chaos underneath. If an attacker gets in, they may take random data, and it will be very hard to assess the real impact on the company.

Coming back to the decision itself, this is also not just one team's responsibility. It is shared across the business. Of course IT and risk officers are involved, but the CEO also needs to understand the main risks, and the board needs to be aware of what can happen and what the possible responses are.

Daniela also mentioned crisis management. You cannot simply find your way out of a crisis if you have never practiced for it.

So yes, this new technology makes things more complicated. But perhaps we can also use the technology itself to learn better how to deal with incidents.

Bruno: Going back to fintechs for a moment: we discussed earlier that fintechs can see what is happening as both a threat and an opportunity. Do you also see them as being able to adapt more quickly than traditional institutions? What is your view, Daniela?

Daniela: Yes, I do.

If we think about how fintechs were designed from the start, many of them were built on cloud infrastructure while everyone else was still operating on-premise. From the beginning, they had to deal with being more exposed to vulnerabilities while also guaranteeing uptime — 99.999 percent in some cases.

So they built redundancy into their infrastructure, with backups and always-on solutions. They also built their organizations in a very agile way, usually with flatter structures where the lines between operational capability and IT capability were blurred.

That created cross-functional expertise at all levels and across all functions, which gave them a degree of resilience that is probably different from traditional players. I would not say higher in every case, but definitely different.

Traditional institutions often have stronger silos and thicker walls between IT and operations, and that is frequently what slows down their response when an incident occurs.

This is where experience and expertise can be shared. I believe traditional institutions can learn a great deal from fintechs.

At the same time, the divide between fintechs and traditional players is not as strong or as clear as it used to be, because many traditional players have gone through several waves of digital transformation. But more can still be done in that direction, and there is a lot to learn from both sides. That is probably the best way forward for responding to these threats.

Bruno: Yes, that makes sense.

Another topic I want to discuss today is vendor strategy and dependence on hyperscale providers. There is a strong trend in Europe and Switzerland to reduce dependence on hyperscalers and build more local cloud solutions. But at the same time, with projects like Project Glasswing, some companies are gaining access to these models and improving their own capabilities.

Andrei, what is your view? Do you think companies already working with Microsoft, Google, or similar providers should extend those relationships in order to get better protection in the short term, or not?

Andrei: There are really two questions here.

The first is whether, if you have tools that improve reliability and security in this very fast-changing landscape, you should extend them to your wider ecosystem — especially suppliers, partners, and distributors. And the answer is absolutely yes.

From the end-user perspective, people do not really care who got compromised. If your name is attached to a product that was compromised, your reputation will be affected, even if it was your supplier or a delegated sales partner. So yes, you absolutely need to think about extending those tools and requiring your suppliers to meet standards for AI safety and security.

Then there is the second question about Project Glasswing specifically.

As I mentioned, you do not need Glasswing. You do not need Mythos. Mythos is not exceptional. It is a little more convenient and a little faster for beginner users, but as CURL and other independent projects have shown, Mythos did not find vulnerabilities that those independent tools and teams were able to find.

We do not need this kind of tool in Europe. We already have the expertise and the people working on these issues.

What we need to do is cultivate those capabilities, especially when access is denied early on. Right now, there is already a two- to three-month head start with Mythos, but other companies could use that time to close the gap and begin patching.

The real problem is not finding vulnerabilities. We have been doing that very well for a long time, even before these new models. The real problem is patching on time.

From that perspective, putting good practices in place now, testing with maybe less well-known models, and coordinating the models in the right way is far more valuable than waiting for a magical tool that will automatically find everything and patch it. That is not how security works.

So my view is: yes, absolutely extend security to the wider ecosystem, but do not wait for the big shiny solution to arrive. There are probably much better things you can start doing earlier.

SECTION 4 — ACCOUNTABILITY

“Regulators Are Watching: Is Accountability Keeping Pace with the Threat?”

Bruno: Makes sense. Thank you for your feedback. Let us now move to the accountability part.

We know that financial institutions are heavily regulated, and regulators are watching closely what firms are doing. We have also recently seen frameworks that have been put into practice, including DORA, which came into force in January last year.

Daniela, how do you see DORA helping with the trend we are seeing now? How can DORA help protect against these threats?

Daniela: I would start by saying that technology moves faster than the law.

By the time legislators come together — especially at European level, where the process tends to be quite slow — and regulate a particular topic, especially one related to technology, the technology has already moved further ahead. So there is always a structural lag.

What must be appreciated about DORA, however, is that it is outcome-based rather than prescriptive. That allows it to remain valid regardless of technological progress.

For example, DORA does not explicitly mention AI or AI-driven cybersecurity threats. But what DORA says is: show that you are resilient, show that you are

prepared for an attack, and do not hide behind your supply chain and then try to shift the blame up or down when a disruption occurs.

Show that you are doing everything you can to prevent it, and then if something does happen — because it is inevitable — show that you have a response plan in place, that it has been tested, and that it is sufficient to match the risk you may be facing.

So it is not prescriptive in that sense. There are some documentation requirements, but it does not prescribe the specific initiatives individual firms must take. That makes it future-proof.

Of course, there are always limitations, but I do not think we need to go back to the drawing board and design DORA 2.0 any time soon just because new technologies are emerging.

At the same time, I still believe legislation is needed, because there are institutions that would not do enough unless they were required to do so by law.

So it is an endless debate. People who are less familiar with the European framework often complain that there is too much legislation on this side of the ocean. But that is what happens when you are trying to coordinate a supranational entity of more than 27 countries.

Bruno: Thank you. To continue with that line of thought, could DORA support, push, or even control third-party vendors and help identify whether they are sufficiently ready and resilient?

Caroline, what is your view? Can DORA help us have better control over critical vendors?

Caroline: I think DORA suggests some reactions that companies can take and provides a framework, but institutions also need to do their part. They need to implement the right processes to make sure they know who they are working with, how they are working with them, and how they can rely on those providers.

I wanted to mention two things.

First, financial institutions need to think carefully about the service-level agreements they have with their providers, including appropriate liability clauses for the specific tools involved. I think that is very important.

For example, if there is a misclassification of a high-risk loan as a low-risk loan, that has an impact on the institution. We cannot simply rely on the provider saying, “That is not our fault.” The institution needs to be able to go back to the provider and say, “You sold us something that did not perform as promised.”

Or, if a tool is used for KYC procedures and customer assessment, then the question is whether it has been tested for fairness and robustness. If that is not the case, the institution should be able to go back to the provider and say, “This is your responsibility; please reimburse the costs, or we are claiming damages.”

So SLAs are important, and they need to be adapted to these new tools and new ways of dealing with such solutions.

The second point is the need for full visibility of the AI supply chain, including the different dependencies across vendors, third parties, sub-partners, and sub-providers. Are we talking about open source or proprietary tools? Are we talking about LLMs? What is actually inside?

Sometimes it is simply described as a black box, and that is likely to create a lot of discussion between the legal team and the IT team, with legal asking what is inside and IT saying they do not know, or think it is fine, or that it is a black box.

That discussion will probably occupy teams for some time, but it is very important, because it also affects intellectual property, data processing, and of course the outcome delivered by the technology.

Bruno: Thank you. To finish this topic on accountability, I want to ask you, Andrei: it sounds as if we are all in the same boat.

How can central and national entities collaborate, share information, or help one another in this situation? What is your view?

Andrei: The thing is that there are already many institutions and initiatives focused on cybersecurity best practices across Europe and Switzerland.

In Switzerland, the National Cyber Security Centre recently signed an agreement with the Swiss Financial Sector Cyber Security Centre to improve collaboration. In Europe, there are also different initiatives that have existed for more than five years, especially around cyber threat intelligence sharing and coordinated red teaming. There is also the G7 cyber expert group, which focuses heavily on these topics.

So we already have very good institutional structures at a high level.

The problem is that the velocity of change, and the speed at which new threats emerge and disappear, has increased dramatically. The pace at which these large institutions operate is simply not fast enough to keep up.

One of the things we face in the field, where people are really working at ground level and often in the trenches, is that sometimes when we try to report findings that could have an immediate impact within one, two, or three weeks, we are unable to reach the right people. Or we are told that we need to go through the official channel, which may only succeed in six to 12 months. That is no longer acceptable.

We need flatter networks. They already exist to some extent in informal form in Switzerland, but they need to be woven much more tightly.

There needs to be more direct coordination between cybersecurity teams across institutions, with no need to get director approval just to speak to someone else. There also needs to be an implicit agreement that this information cannot be used negatively or offensively.

In my opinion, shifting things further down, much closer to the trenches where threats are first seen and felt, is one of the biggest challenges and one of the biggest changes that needs to happen quickly.

SECTION 5 — WHAT NOW

"90 Days to Act: Priorities, Investments and the Honest Conversation"

Bruno: Thank you. Okay, we are now moving to the next topic. If you have any questions, please write them in the chat and we will come back to them at the end.

So now the question is: what do we do now, in practical terms? I want to ask Daniela a quick question. If a client comes to a financial institution and asks, "Are you protected?", what should the answer be? And what are the main actions we need to take urgently to address that?

Daniela: I think there is only one correct answer to that question, which is: "I don't know, but I am prepared."

There is no such thing as being able to say, "I am 100% protected against any type of attack." We simply do not know. There are unknown unknowns, and that is the new paradigm and the environment we live in.

There are uncertainties, and what is required is the ability to respond to those uncertainties. That is incident readiness, as Caroline mentioned earlier. It is about complementing the protection that already has to be in place.

All the traditional — or rather, standard — security measures we have been discussing for years must be in place. And, as the UK AI Security Institute found in its assessment of Mythos, the same basic security methods still need to apply. As Andrei said at the beginning of this webinar, the standard methods of security need to be in place, and then more needs to be done to expedite the response and make it as agile as possible, so the organization can respond from all angles an attack could come from.

No one can honestly say, "I am 100% protected." We should not be skeptical of someone who says they are not 100% protected, because that is the truth. Instead, we should question those who underestimate the level of threat.

Bruno: Perfect. Thank you. Let's now have a final question for all of you.

At the beginning, we discussed the key pillars for financial institutions: trust, stability, data protection, and data security. Now that we know what is happening, I would like to close this exchange with one commitment, one action, or one message from each of you to share with the audience.

Let's start with Andrei. What would be your main takeaway?

Andrei: Start with proper cybersecurity basics.

Do the things you may have been putting off because you thought you were too small or because you believed attackers had more important priorities. With automation, attackers no longer care how small you are. If you are profitable, they will go after you.

You need to start doing the basics well now.

Bruno: Great. Caroline, over to you.

Caroline: I would go back to regulation and remind everyone that regulation exists to protect customers, citizens, and their rights.

We cannot neglect those rights for the sake of profit, speed, or efficiency. At the end of the day, we are protecting customers and citizens, and it is very important to set the right measures and limits for that purpose.

Bruno: Thank you. Daniela, over to you.

Daniela: For me, it is about strategic response.

Regulators no longer care about how much is spent. They want to know how institutions protect themselves from cyber threats. They care about how the money is spent, and in fact, what they really want to see are the outcomes.

Show exactly how your organization is able to respond. Show that there is awareness at all levels of the organization. This is no longer just a conversation for the IT department. There has to be recognition that these threats can come from any direction.

Let me leave you with a practical example from the UK. One of the largest companies, which also has a strong financial arm, was recently attacked through a phishing attempt aimed at middle management. It was a fake email, but it looked legitimate and asked for specific information, which the hackers then used to bypass security and steal customer information, credit card data, and so on.

The attack did not come through the IT department. It came through middle management. That shows that the operational setup was flawed and that the scenario had never been practiced.

So my invitation to everyone is this: practice these scenarios, learn from other people's mistakes, and make sure you are well positioned for any eventuality, even if it seems low probability.

Bruno: Thank you, all of you, for these inputs.

Let's move to the chat and the first question we received. Adrian asks: if offensive systems are ahead and defenders do not yet have equivalent tools, what defensive counterpart should we be trying to build? Andrei, what is your view?

Andrei: I think there are several elements to defense.

The first is automated patching. Right now, when a vulnerability is found and you need to fix your code, you can use automated tools to rewrite it, but the problem is that they often slightly change the logic. Often enough, in fact, that you cannot fully trust them and still need a human to review the result, which takes time.

We need to invest more effort in making sure we can deterministically verify that everything has gone right.

The second element is defensive tools. There is a lot of discussion about deploying defensive agents, but the problem is that they can become a kind of paperclip-optimization machine. The safest system is one that is fully online. The only database that cannot be leaked is the one that has been deleted. So if your database gets deleted by your own defensive agents, that rather defeats the purpose.

A lot of research is still needed to ensure security while maintaining availability, and we are not there yet. That is the second area.

And finally, we need to make sure these rapid-response defensive agents are actually understood, given permissions that are clear, and can be stopped if they start causing problems.

There are companies and entities claiming they have tools to manage permissions through large configuration files, but they are not really ready for real-world production, where someone needs to review and decide which permissions should or should not be granted.

So we need innovation in this direction. There are some promising leads, but nothing ready for production yet.

Bruno: Okay, perfect. Let's see if any other questions come in over the next minute or two.

In the meantime, Daniela, a quick question for you about the regulatory trend. Do you see new regulation coming, or regulators publishing something soon in response to these new models? What is your view?

Daniela: Yes. In my experience, what has changed quite a lot is the approach regulators have been taking across Europe.

Traditionally, the approach was more prescriptive. Now we have evolved toward a more outcome-based approach, although there are still nuances. For example, in Germany the approach is still quite prescriptive. It is very clear what financial institutions are expected to do, including timelines for incident response, resolution, and continuity.

In other countries, the approach is becoming more and more outcome-based.

There was a point Andrei made earlier about informal networks for information sharing. I see some countries moving in that direction, particularly in the Baltics. It is easier for them because they are smaller countries, they probably regulate fewer financial institutions, and those institutions tend to be very tech-savvy.

I also see this in Switzerland. FINMA, especially on cybersecurity matters, seems to me to be much more responsive than it is on some other, more traditional financial issues. I do not know whether that reflects an evolution from the past, but it may be a question for the audience or the other panelists.

And finally, I can comment on the UK. The UK has always followed a more Anglo-Saxon principle, requiring firms to declare their own impact tolerance and then allowing the regulator to assess each firm against its own declaration. You never know exactly what the regulator is looking for, which can be confusing for some, but it also allows more freedom and avoids a one-size-fits-all approach.

If this is your business model, show me that you can manage your business model. If this is the approach you want to adopt, show me that you have the customer in mind.

That is also Caroline's point: the outcome matters. The protection of the customer and the trust in the system are the two principles that are common to all regulators across Europe, and possibly the US as well.

Bruno: Thank you, Daniela. And thank you all for your contributions. This has been a really interesting discussion and interaction.