BEYOND blue

NEWSLETTER
JUNE 2022

As the EU Commission's draft Digital Operational Resilience Act (DORA) consultation approaches finalisation, we take a look at what DORA means for firms and consider how it relates to other resilience regulations such as the v/FCA's Operational Resilience Framework. We also reflect on the emergence of cybercriminal gangs such as Lapsus$, consider what constitutes best practice in API security, and discuss what's next for Operational Resilience implementation programmes.

David Ferbrache, OBE

DORA THE IMPLORER

In September of 2020, the EU Commission published its Digital Finance consultations package (DFP). This set of proposals was designed to help foster an EU-wide strategy for both maximising the opportunities and minimising the risks associated with digital transformation of the financial services sector (FSS). On the opportunities side, the DFP seeks to outline ways in which the EU financial sector can adjust to enable the provision of innovative new financial products. These include the de-fragmentation of the digital single market through measures like the cross-border harmonisation of digital identifiers, innovation-friendly adjustments to the EU regulatory framework, and a strong focus on new and improved rules on sector-wide data sharing. Meanwhile, the cornerstone of the risk minimisation side of the DFP is the draft Digital Operational Resilience Act (DORA). This ambitious new legislative proposal attempts to expand on existing EU cyber risk management mandates to provide a much more comprehensive framework encompassing both standards and controls. This makes the DORA a major regulative undertaking. One which, if approved in its current form, will require expensive and extensive compliance programmes across the EU financial sector and beyond.

Those familiar with the PRA/FCA's recent Operational Resilience regulatory framework (UKOR) for the UK financial sector might be forgiven for assuming a certain continuity between the UK regulations and what the EU is proposing for the DORA. For one thing, there is the simple fact that both are nominally regulations concerned with the same thing, i.e., the operational resilience of financial institutions. For another, much of the initial commentary on the DORA has largely assumed that it embodies the same central objectives as the UKOR. Namely, that the DORA's central goal is in 'ensuring firms […] are able to maintain resilient operations through a severe operational disruption'.[1] This is a natural assumption to make if one's starting position is the conception of operational resilience as found at the heart of the UKOR. The concept of resilience embodied in there is one on which it is the property of systems which enables them to adapt to surprising, disruptive events to avoid sudden failures. The system in question in the UKOR is a network of business operations provided by financial institutions engineered towards the provision of critical services to customers (the so-called "important business services"). Hence, the UKOR definition of operational resilience is as the capacity of FIs to adapt in the face of major disruptions so as to be able to continue to provide important business services.

[1]Deloitte (2021): https://www2.deloitte.com/lu/en/pages/financial-services/articles/the-eus-digital-operational-resilience-act-for-financial-services-new-rules.html

OPERATIONAL RESILIENCE UKOR

the capacity of FIs to adapt in the face of major disruptions so as to be able to continue to provide important business services

In fact, the superficial parallels between UKOR and the DORA conceal the fact that the concept of resilience at the heart of the DORA is not like this. The DORA's notion of resilience is closer to something like system strength or lack of weakness. To be resilient, on this picture is for a system to manifest a high-degree of tolerance to disruption given its overall robustness and lack of vulnerabilities. Furthermore, the nominal "operations" referred by the DORA are different also. They are not business operations which support external-facing services, but the internal operations of FI-hosted information systems and networks on which digital finance depends. Hence, the DORA's notion of operational resilience is something closer to the strength/lack of weakness of FI-hosted digital infrastructure. For this reason, pursuit of DORA-style operational resilience is effectively identical to the pursuit of a low level of operational cyber risk.

OPERATIONAL RESILIENCE DORA

the strength/lack of weakness of FI-hosted digital infrastructure

These fundamental conceptual differences manifest also in a difference of regulatory approach. The UKOR is a principles-based piece of regulation. It is light on specific mandates for compliance and is at times deliberately non-specific. This is a feature not a bug. The motivation behind adopting this approach is precisely that it enables FIs a high degree of flexibility in interpreting its mandates and experimenting with different approaches to compliance. Given the essential emphasis on dynamic adaptability in the UKOR conception of operational resilience and the vastly different organisational structures and business logics across the UK financial sector, this approach has obvious advantages. Conversely, the DORA is a standards and controls-based regulatory framework. It outlines highly detailed and specific instructions for FIs both on how they are required to configure their networks and on how they are required to manage cyber risk. This too is consistent with the mandate of the DORA. Lowering operational cyber risk will be a standards and controls oriented endeavour.

Importantly, nothing about this divergence between UKOR and the DORA means the two cannot be complementary. Both have essentially different objectives and are driven by different concerns. UK FIs with UKOR implementation programmes have continued to maintain cyber risk management functions within their organisation precisely because lowering cyber risk and shoring up operational resilience are, on the UKOR definition of the latter, discrete tasks. For the same reason, there is no principled basis for thinking that implementing UKOR should serve as an impediment to implementing the DORA.
In fact, there is more likely to be tension between the implementation of the DORA and the EU's Network and Information Security Directive (NIS). For whilst the NIS Directive is principles-based and the DORA draft specifies that its mandates will override those of NIS where conflicts arise, there are likely to be substantial difficulties ahead for those firms subject to member state specific standards and controls legislation designed to ensure compliance with NIS.[2]

Similar problems lie ahead for global firms. These organisations are already required to comply with complex and detailed standards and controls mandates emerging out of non-EU jurisdictions such as India, Singapore, and the State of New York's Department of Financial Services. The DORA will add a burden to these already sizeable efforts, containing as it does obligations on everything from reporting and testing through to incident and third-party risk management. Moreover, given the specificity of many of the requirements contained in the DORA, there is also the very real possibility that regulatory conflicts will emerge. Even in the EU itself, the interaction between DORA and extant regulations such as the Payment Services Directive 2 (PSD2) have yet to be clarified. For example, both require incident reporting but assume different timescales and models.

As the DFP and the DORA edge closer towards finalisation, firms will need to bear the above considerations in mind. For UK based firms with EU operations, it is important to avoid equivocating the two senses of operational resilience at stake in UKOR and the DORA. There are few if any commonalities between the two frameworks, despite their superficial closeness. For the same reason, EU firms should take note that, in implementing the DORA, they are not shoring up their capacity to adapt to disruptions in a way that secures their core customer facing services. Meanwhile, EU-based and global firms should prepare themselves for a significant challenge. Implementing the DORA's extensive set of regulatory requirements within the two-year window required will take careful planning and extensive resourcing. This will be exacerbated further by the complex regulatory conflicts which will likely emerge along the way across different jurisdictions.

[2] DORA (2020), §16, p.16

WASIM AKHTAR ON API SECURITY

Wasim has over 20 years of technical experience in delivering resilient, secure IT architectures. He has a particular interest and specialism in virtualisation, Cloud, and networking technologies. In a previous life, he spearheaded the architecting of a multinational IT network for a large aerospace manufacturing firm.

In cyberspace we tend to have retrospective views on our digital world. It only seems to be after the fact that we wish we had done more than just utilising the security controls bundled into our various product licenses and subscriptions. This is partly due to the fast pace of digital innovation. We often now adopt new technologies whose security landscape hasn't yet fully evolved. The pandemic has certainly accelerated this trend and, in some cases, clashing business priorities have meant that security has become an afterthought.

A good example of this concerns the use of APIs. These are now commonplace in our everyday lives. They make possible our digital navigation through cities, payments for products with a tap, and even the summoning of a car. They are also near essential to contemporary Industrial Control Systems and are used widely across Critical National Infrastructure (CNI).

Inevitably, as well as making our lives more efficient, APIs also introduce new vulnerabilities. Increasingly, a variety of actors (both malicious and non-malicious) are proving they can manipulate API architecture for monetary gain. Targeted behavioural advertising, price manipulation, scalping of rare goods and anti-competitive practices in e-sports are increasingly common. There are also sophisticated nation-state and criminal actors using these vulnerabilities to perpetrate crimes against the financial sector, election fraud and attacks on CNI. The techniques used often combine the power of automation with the leverage of common API vulnerabilities to mobilise techniques such as credential stuffing, account takeovers and session hijacks.

Despite the increasing importance of API security, there remains a general lack of awareness about how serious and ubiquitous the associated vulnerabilities are. This raises the following question: "What should our approach to API security be?"

Here are some general tactical pointers which can help you and your organisation down the right path.

Don't just rely on a Web Application Firewall (WAF)

Whilst it may provide you with regulatory compliance, a WAF is just a rules-based system best suited for legacy apps. Crucially, it is unable to look at the context of the application which is key when considering API security.

Don't just rely on code obfuscation

Code obfuscation is commonly used as a first line of defence against reverse engineering attempts by hackers. We need more than this. Crucially, we must have mechanisms able to detect and control who is accessing our API server. Mobile App Attestation solutions are useful when seeking to authenticate the application and user, not just the code.

Don't just rely on security-as-a-feature

Enabling security features on your Content Delivery Network may not be the best fit for the purpose of protecting APIs. API specific solutions allow additional features such as API call tampering prevention, separating authorisation from the service and hiding these flows from the client.

Consider your API lifecycle and add security layers to each phase

Use a forward looking, purpose-built security product from a specialist vendor at each level. This defence-in-depth approach enables you to provide context and to get closer to the workings of your application.

Code Compile phase

Hardening code is a technique to test software against abuse cases that could be encountered in a hostile environment. When code review is performed, it is important to evaluate the impact of a software defect because it may result in a real vulnerability that will represent an exploitable weakness.

Mobilise passive scanning

After an application is deployed, a passive scanner looks at the requests and responses and analyses ingest traffic. It can also act on downstream devices. This is good for finding problems like missing security headers or missing tokens. Note, however, that it will not help find vulnerabilities which require malicious requests to be sent – that's the job of the active scanner.

Mobilise active scanning

Active defence is a solution that sits active within the API. It enables one to examine traffic and take actions. Most importantly, if it spots anything awry, it can deny access to the API gateway completely. This is a technique that gets us up close and personal to the application and adds context.

Use "fuzzing"

Scan regularly for vulnerabilities and determine how your service behaves when getting random or unexpected inputs. Fuzz testing sets operation parameters to unexpected values to cause unexpected behaviour and errors in the API backend. This helps you discover bugs and potential security issues that other QA processes may miss. Implementing modern fuzzing into your CI/CD pipelines will enable you to build a reliable API endpoint testing that provides code coverage visibility and simplifies your debugging efforts.

SECURITY SPOTLIGHT – LAPSUS$ UPS THE ANTE

Since the start of 2022, the Lapsus$ hacking group has become one of the most successful and feared cybercriminal gangs in the world. In just four months, the group have been responsible for high-profile hacks of T-Mobile, Globant, Okta, Ubisoft, Samsung, Nvidia, and Vodafone. Perhaps most spectacularly of all, Lapsus$ hackers managed to break into Microsoft's Azure DevOps server in March and were able to obtain source code from three of the company's services: Bing, Bing Maps, and Cortana. According to Microsoft, the attackers gained access to the network using a single Microsoft employee account. Using this, they were also able to access Microsoft's servers and steal the company's source code. Soon after the hack, the criminals boasted about it on their Telegram channel. Lapsus$ released a torrent for a 9GB zipped bundle containing the source code for over 250 Microsoft-related projects. They claimed that the torrent file contained 90% of the source code for Bing and around 45% of the code for Bing Maps and Cortana.

To gain initial access, the group tend to focus on compromising user identities and accounts. This is often achieved by either acquiring credentials from sites on the Dark Web, scanning public repositories for exposed credentials, use of the Redline password stealer, or recruiting personnel at targeted firms to operate as insiders. They then tend to use publicly available tools to search through an organization's user accounts for employees with higher privileges and broader access. Often this will be followed by theft of further credentials from development and collaboration platforms like Jira, Slack, and Microsoft Teams. According to Microsoft's account, Lapsus$'s attack on their network followed exactly this pattern. First, they used a SIM swap attack to gain possession of an employee's phone number and text messages. This then enabled them to gain access to multi-factor authentication (MFA) codes needed to log in to a company. They then used these credentials to access source code repositories on GitLab, GitHub, and Azure DevOps.

The brazenness and ambition characteristic of this attack is chastening. Mercifully, however, the fortunes of the group appear to have taken an equally sudden turn for the worse in recent weeks. Towards the end of March, a group of seven UK teenagers were arrested under suspicion of leading the group by the Metropolitan Police. One in particular, a 16 year-old from Oxford, is suspected of being one of the leading public-facing figures within the group, known by the handles "White" and "Breachbase". Quite what this means for the future fortunes of the group – who are believed to be primarily based out of South America – is unclear. However, it would be foolish for security professionals to take their eyes off the road. Unlike many other hacker groups which comprise a closed network of a small number of members, Lapsus$ represents a quite new kind of threat: the open-source hacking collective. Part of the reason for their remarkable recent success is that Lapsus$ makes use of social media platforms like Telegram to drop data useful to other hackers in a public way. This enables potentially thousands of hackers from around the world to contribute to their efforts. Tackling an agile threat on this scale is uncharted territory for the security industry, but one which it must prepare itself for. If the beginning of 2022 is anything to go by, this new trend could prove to be highly destructive and harder to defend against.

OPERATIONAL RESILIENCE – SELF-ASSESSMENTS

The first Operational Resilience deadline has been and gone. With its passing, the sector has breathed a collective sigh of relief. The first version of the Self-Assessment has been approved by firm's Board. But how difficult was it to obtain that approval? What questions did the members of the Board ask? How will these questions (and the attempts to answer them) shape Operational Resilience programmes in the lead up to the next deadline of March 2025?

It is likely that Boards and the role of SMF24 will have quickly turned their attentions to the implications of the Self-Assessment. The remediation of vulnerabilities unearthed through mapping of critical resources to important business services (IBS) and subsequent scenario testing will come with a significant price tag for many. The crucial question they will be pondering is the following: "When only considering business impact and not the likelihood of the scenario, how should a firm seek to prioritise remediation?" Few organisations are working with a bottomless pot of funding and so some mechanism of prioritising is required. Will it be dependent on the number of IBS impacted by the vulnerability, the severity of the scenario that identified the vulnerability or the cost and complexity of the remediation? Whatever blend of factors that organisations consider, firms have to be comfortable in their justification if they choose not to remediate. The regulator will expect Boards and ultimately, the individual in the role of SMF24 to be able to stand behind the statement: "We are resilient enough".