What We Call "Security"

Blog series by [Greg van der Gaast]

brought to you by
[VM2020 Solutions] and [Hitachi Vantara]

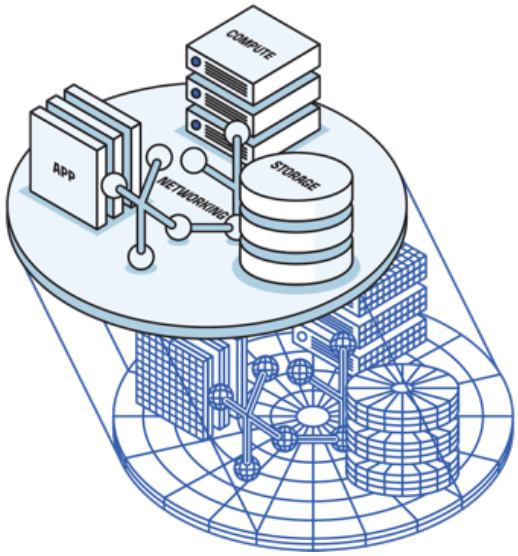

Table of Contents

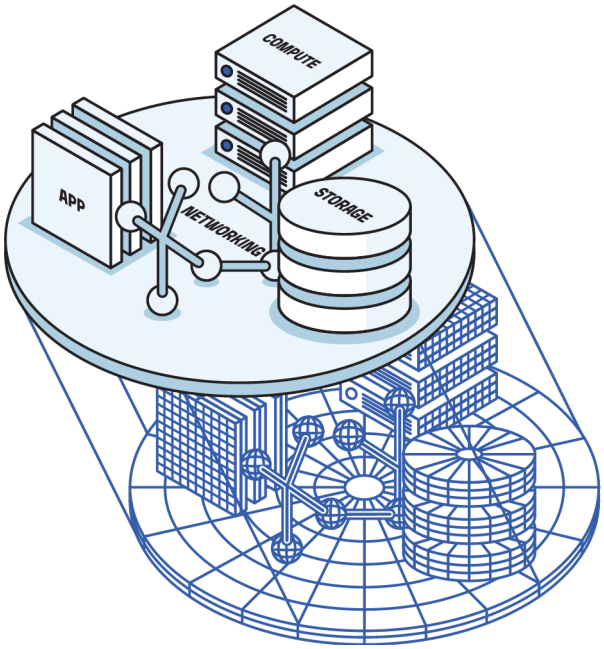

To learn more about Thin Digital Twins technology,
or get a free copy of the book "What We Call Security", contact us at *info@vm2020.com*

Blog (0/6) – Introduction

Security, Strategy, Storage, and Recovery

Welcome to this new blog series in which I will be offering a different look at storage and recovery: One from a strategic *security* point of view.

The goal is not just to better explain the relevance of storage and recovery to current and aspiring Information Security leaders, but also to introduce some concepts around "Security as Quality", what I call "Inherent Resilience", and presenting a different approach to Risk Management to help practitioners move things forward.

The reasons being that, in order to understand and appreciate the value of new approaches, it's important to understand the bigger picture and how things fit together.

I hope these insights also help IT Operations teams and CIOs realise additional security potential from platforms and processes they may not have thought as relevant.

So, what is "Inherent Resilience"? Well, it's a term I've coined to define our ability to not get knocked down in the first place. In other words: Not reaching the point of needing recovery in the first place.

I've phrased it this way because in the security field "Resilience" has come to mean something closer to *recovery* while I think our long-term focus should be focused more on not getting breached in the first place.

And if you're confused that a blog about storage and recovery is talking about avoiding the need to perform recovery in the first place, then hold on to your hat because we are going to discuss how to use your recovery capabilities to improve your chances of never needing to use your recovery capabilities. Intrigued? Stay tuned.

But first, an introduction. My name is Greg van der Gaast. I have 25 years of experience in information security. My first job was offered to me when I was 17 and some federal agents from the FBI and US Defence Department decided to make a house call. I won't bore you with the details as to why, but it involved the NSA, CIA, DIA, and some nuclear weapons. I have seen hundreds of breaches and caused a few myself - once having been labelled as one of the 5 most notorious hackers in the world.

One of my biggest observations after switching sides, was why my job as the attacker was so easy, and why what most of the security industry was doing wasn't making it much harder. As a result, I've always had what people consider a "maverick" approach to security because I believe in doing what works, long-term, sustainably, rather than the status quo.

Over the last 15 years I've since built security programmes for companies ranging from hot start-ups to Fortune 500's, lectured at universities on security strategy and leadership, advised cyber-insurance

companies on due diligence (hint: quality of business processes is a far better indicator of risk than the presence of security controls), and currently assist security vendors in helping their customers get more value from their offerings.

A quick note: I am *not* a storage and recovery expert, but I want to take you on a journey to look at our security challenges, how I've tackled them, and, finally, what role storage and recovery has played for me as part of that bigger picture.

I also want to add that Hitachi Vantara, while sponsoring this series, has given me no further instructions on what to write other than helping the community at large. These are my thoughts and experiences, and I thank them for allowing me to share them with you.

But before we get started on how to develop *successful* security strategies and approaches and where storage and recovery fit in, we must look at the overall trends in our industry to see how things are going:

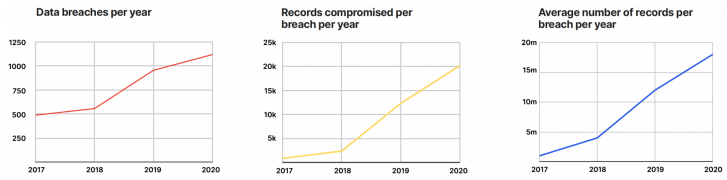

Source: Imperva

Not great, are they?

And this worrying trend is happening *despite* ever-increasing spending on security:

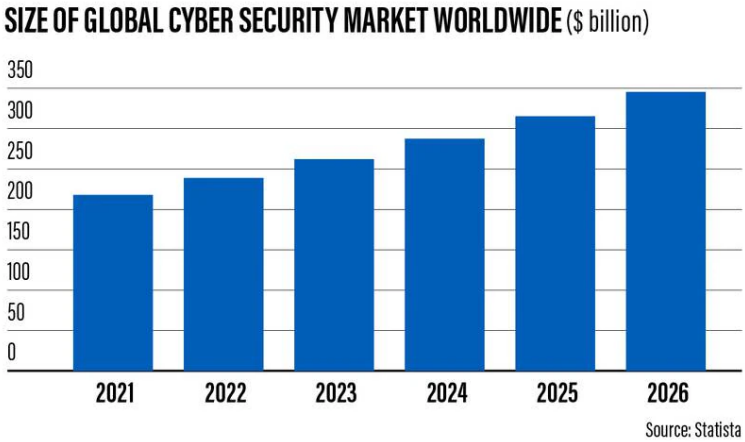

In most situations, certainly at this scale and over a period now spanning well more than a decade, investment in an overall approach is expected to have an impact in reversing or reducing what it is it's trying to combat.

Instead, many practitioners, vendors, and experts use these figures to point out how prevalent and sophisticated attackers are, and that we must therefore double down further on our investment in "cyber."

I have instead come to see it as a rather damning indictment of just how poor our current approach is. Dozens of times I have seen companies with millions spent on security, with NIST and ISO frameworks in place, and still be absolutely trounced by a bored teenager with a laptop.

If we had an effective and sustainable approach to reducing issues, we should be seeing trends like those in mature industries where they too are fighting to reduce the number of incidents. They identify root causes *no matter where they are* and address them, upstream, pre-emptively.

Take the aviation sector for example, which addresses issues as far upstream as possible regardless of how distant they seem from most people's concept of "aviation" (Everything from engineering, metallurgy, corporate culture, drug use, the tone of alarm sounds, control ergonomics, human factors, etc.) to drive a reduction in possible failure points:

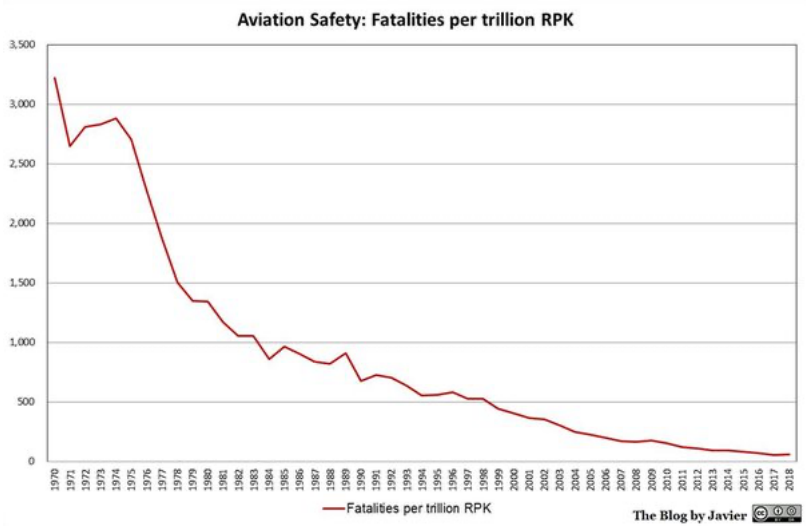

That is what the results of an *effective* approach look like.

If we found that the bolts holding the wings onto a plane's fuselage were coming loose during flights, we wouldn't set up a function where we employ thousands of people to retorque bolts after every flight, forever. We'd make a change to the design or manufacturing process *once*, then remediate what was in the field. And yet, most security work has more in common with the former than the latter.

So, what should we be doing in Information Security? It seems clear to me that a change in approach is needed. But what? And more importantly, based on what principles?

Let's begin with some opening questions:

Have you ever considered…

- That security vulnerabilities are *defects*? Whether it be in code, architecture, design, maintenance, process, or even human behaviour?
- That, to rectify this, Security might ultimately be more effective as a business *quality* function?
- Why we tend to focus on threats and protecting vulnerable applications, systems, and infrastructure, rather than on changing the business processes that lead to their vulnerability in the first place?
- If we could drive improvements to security without having to continuously (and unsustainably) increase the scale of security operations?
- How mature industries like automotive, manufacturing, or aviation stop issues from recurring or occurring at all? And how some of approaches could be relevant to Information Security?
- Whether Risk Management could be simpler if we calculated backwards from business downtime, rather that the innumerable arbitrary compounding variables that might lead to that downtime?
- How storage and recovery capabilities can allow us to shift more resource towards a strategic security approach, rather than mitigation and firefighting?
- How recovery should be implemented to ensure reliable recovery if things do go wrong?

Over the course of this six-part series, we'll be exploring all these questions and what they mean to improving the security posture of our organisations. We look forward to having you join us for these insights. Don't miss it!

Blog (1/6) – Security as Quality

We ended our introduction to this series with several questions. Questions I've found myself asking over my 25 years working in Information Security. Some of these questions stemmed from others, and some led to more questions, but they did take me to a general conclusion:

The way that we currently practice security, overall, isn't just ineffective as shown by virtually every statistic. It's also *weird*.

We have developed very niche approaches to solving problems in our industry. But if you applied analogous practices in other industries, we would likely leave people scratching their heads at what we were doing. Conversely, other industries have developed and refined approaches that deal with analogous problems very well, yet we reject them as not applicable.

I've seen a particular Sun Tzu quote used in several security presentations talking about the importance of threat intelligence, detection, response, logging, monitoring, and a whole host of other security solutions and practices.

Know thy enemy.
Know thy self.
A thousand battles.
A thousand victories.

Most of their narratives highlighted the importance of "know thy enemy", but in my eyes failed to realise that they were omitting the *know thy self* part.

In short, I believe we are too quick to focus on the threat actors while we ignore what makes them *threatening* to begin with: Our own [excessive] vulnerability, or vulnerabilities.

But what is a security vulnerability? You could argue in its simplest form that it's a defect, a quality defect. A defect in code, in architecture, in a workflow that doesn't allow maintenance, in a procurement process that brought in something else that had defects, in any business process really. And yes, every IT process, no matter how small, is a business process to some degree.

Would we allow planes to leave the assembly line with known defects that could make them crash, and then leave it up to some "operations centre" to try and detect when these defects start causing dangerous situations, in the hope of stopping the worst-case scenarios?

The "operations centre" may play a role in detecting and trying to mitigate such an issue, but regardless of the outcome we would immediately address it in the design of the aircraft, on the assembly line, the very first time it happened, to make sure the risk is never introduced again.

At the end of most automotive assembly lines there is a person doing quality assurance. Their job is to see if any defects have been introduced in any of the 100+ sequential stations in the assembly line building the car. If they do find a defect, like the steering wheel being mounted off from centre, the car goes back to that station to be fixed. If this defect is significant or occurs more than once, then the time is taken to examine the cause and re-engineer the process *at the station where it was first installed* to prevent it from happening again.

They do not hire an ever-increasing army of quality assurance people to fix defects at the end of the line, nor would they, if a defect kept recurring, just keep sending every affected car back to the line without addressing the process that caused it to have that defect in the first place.

But these rather common-sense approaches to quality management seem to have largely escaped the field of Information Security. As a result, the more our organisations grow, the more technology we take on, the more security issues result, and the more we constantly need to firefight. I would argue that much of what we have come to call "security" has more in common with firefighting than "securing".

We are not only failing to learn from what other industries have long ago figured out, but we are also missing out on the economies of scale of putting quality into the build process (of code, systems, architecture, process, etc.) rather than endlessly rectifying things at the end. Plus, we've reached a point there the volume of firefighting is so high that we can't keep up despite forever increasing our investment in "Security".

Security is incredibly well placed and effective at detecting these quality issues, but we rarely leverage the information to actually fix the processes that cause them.

It's the reason why both incident and spending numbers keep climbing. It's the reason we too easily now say "it's no longer a matter of if but when" when most breaches have been shown to be readily preventable.

You might recall the trends below from this series' opener.

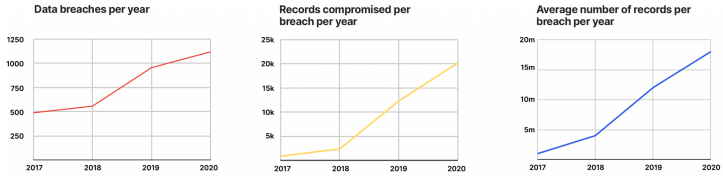

Source: Imperva

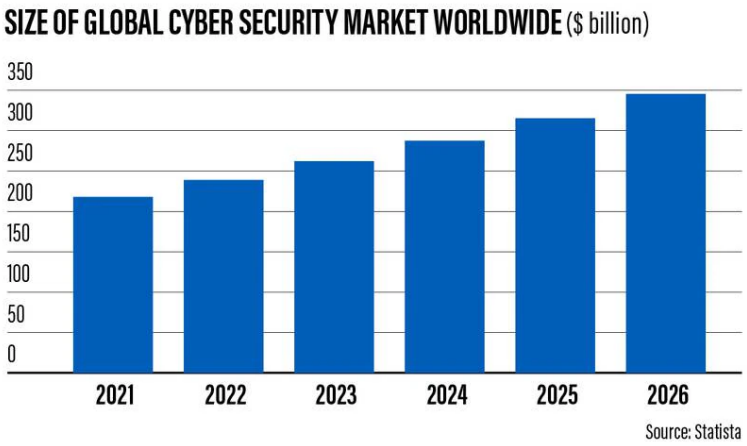

Too often, security practitioners' minds focus immediately on "compensating controls" and mitigation efforts without thinking about why they are needed in the first place.

Worse still, just like every other system, the effectiveness of security controls is also dependent on the integrity of dependencies which are often not considered. For example, your IAM service validating your users cannot be trusted if it has a code defect allowing the injection or modification of accounts, or running on a server missing critical patches that allow someone to compromise it and have full control over your IAM setup.

Incidentally, this is the fundamental reason I tell cyber insurance providers doing due diligence to look at the maturity of business processes (something often primarily driven by company culture) rather than the mere presence of security controls. Without good process, the security controls themselves are likely to be ineffective or potentially compromised.

Another rarely considered and counter-intuitive angle: The needed presence of many compensating controls can point to poor business process, and higher risk, in of itself. Why were so many compensating controls needed? Why wasn't the security built in earlier? What internal issues led to that?

If a car company built a dedicated team of people to find and realign wrongly installed steering wheels at the end of its assembly line, but neither the company nor that team had bothered finding ways to make sure they were mounted correctly in the first place, wouldn't you have some questions?

Questions about the process, about the company, about the people there?

Wouldn't we be asking ourselves why they hadn't focused on the much easier solution of maybe creating a jig or a template when first installing the steering wheels so that all this work at the end wouldn't be needed? Wouldn't that be faster, cheaper, and generally more beneficial to the business?

This is my personal definition of *strategic* security. Asking the questions that lead to approaches and solutions that create long-term reductions in issues and put the business first, rather than building up

operational work (with no long-term benefit) to mitigate ever-increasing symptoms of a more fundamental problem.

So, what does this have to do with storage? Or even recovery?

In some ways, not much. In other ways, everything.

This strategy, or philosophy, or *common sense,* of Security as Quality shapes how I and some others approach our roles as CISOs. It's how we strive to reduce attack surface over time, make sustainable progress rather than firefight, and lower costs. It's also a way of doing things that offers numerous additional business benefits to our organisations because improving quality improves much more than just security. Storage and recovery play a crucial role here in two ways:

The first reason is simply the security status quo that we usually start with. The common industry practices have created a "way of working" where an organisation's vulnerability is ever increasing in scale and complexity. To the point that, using present-day security approaches and solutions, that vulnerability can no longer be effectively managed with the means at our disposal. This can be seen by the trends of ever-increasing quantity and scale of breaches despite increased spending on security.

We have created a space where it's very much a matter of "If, not when." And while this phrase is technically true, in the same way that an asteroid will likely eventually hit the very spot I'm sitting in as I write this, I fear we've come to use it as an excuse to not do better. Or we've just failed to realise there are other ways of approaching the problem.

As a result, recovery has become of paramount importance and grows more important by the day because, as the trends demonstrate, in generalised terms at least, our current approach cannot stop the breaches from coming.

The second reason is however far more compelling to me: Recovery capabilities give us a safety net that can empower us to change things for the better. It can allow us to deprioritise firefighting and instead prioritise a Security as Quality approach that drives lasting improvements.

From a pure Risk Management perspective, recovery allows us to limit the maximum impact of a breach. This means we can (and should, based on risk management logic) reallocate more resource to where it can drive greater change.

We can refocus to reducing how much risk our business processes introduce, driving long-term reductions to how much Security needs to manage, in turn freeing up evermore resource for more proactive work. This exponential growth in our power to affect change is what can turn the tide and make our organisations increasingly *invulnerable*.

That brings me to the end of this instalment which I hope has fuelled some thoughts on how we can approach security differently. In our next instalment I'd like to share how I structure security

programmed and how recovery capabilities give us the freedom to reallocate resource away from reactive efforts towards ones creating long term improvements, all while helping us sleep at night too.

Blog (2/6) – Building a Programme

In this instalment I want to share with you how I structure a strategic security programme as a CISO. Both in the hopes that it can help you drive change in your own organisation(s), but also to help the practical understanding of some of the things we have discussed so far on in this series both in terms of how and why.

This will be a high-level overview at best, but something we will explore in detail in an upcoming book graciously sponsored by Hitachi Vantara.

My definition of a "strategic" or "business-level" security programme is similar to any other security or quality programme. It involves defining the ideal state of things, how we maintain that state, and capturing the necessary support and authority to do it. The real difference is the scope at which we address the problem:

Most security function and programmes focus on having the capacity to deal with all the issues the current business processes have and will produce, forever.

Conversely, I prefer to focus *primarily* on influencing business process throughout the whole organisation to continuously reduce the number of bad outputs (vulnerability, attack surface, obstacles to security maintenance activities like patching, etc.) they produce in the first place.

The first is building operational security capacity, the second is shaping the business so that things are inherently more secure and less and less operational security capacity is required over time.

So, what does such a programme look like? Well, it can take many shapes. I've evolved, sometimes radically, how I approach this over time to not just better reflect each organisation, but also to incorporate new knowledge and ideas. I hope that, should you decide to take some of the concepts and structures presented here, you also adapt and evolve them to work even better for yourself and others.

The first step is very simple: I create a framework.

I start by drawing a big square (more a vertical rectangle for me) that acts as both a box and a wireframe.

The purpose is simple: it holds everything. It's a container that everything will live in, and that will connect them all together. This overcomes the first problem I often see in organisations: discombobulated documents and policies, rarely in homogeneous formats, scattered all over the place, and with no full inventory anywhere.

Inside this box I create layers, or domains (I'll used "domains" going forward). And within these are a variety of documents. These documents can be standards, declarations, processes, policies, definitions, anything really. These are what define the desired states of things; of the organisation, of

our own efforts and mandate, of our needed authority, and of the approaches we'll employ to deliver on our mission.

How many domains or layers there are can vary for each organisation, and as mentioned you should tailor this to what works best for you. But as an example, the last framework I developed for a SaaS B2B services company contained the following domains:

1. Executive
2. Programme Definition
3. HR & Legal Integration
4. IT Operations
5. SaaS & Business Applications
6. Product and Engineering
7. Human Factors
8. SecOps
9. Compliance
10. Commercial
11. Business Stubs

1. Executive: This always includes an Executive Charter that defines hard rules for security scope, responsibility, and authority, signed by the senior management team (The perfect place to pitch the security as quality concept and its many non-IT business advantages). As well as a written Security Strategy (business-centric, so well beyond IT) that explains and justifies the approach (including the framework), team structure, timelines, etc. involved in delivering it.

   This is critical. You need a business security strategy, and you need real executive support. Many organisations have neither.

2. Programme: An overview of the whole programme, inventory of its components, continuous improvement process, how the activities will be scheduled and tracked.

3. HR & Legal Integration: I need to define Acceptable Use Policies enforced by HR, I need defined roles and associated access, I need HR integration for automated/effective JML, I need to work with HR and sometimes legal on investigations. All this is defined here. I also need integration with Legal for contractual reviews, certain incidents, ensuring contractual requirements are captured in IT and security service delivery, etc.

4. IT Operations: Define how we do everything IT related, from provisioning users, asset management, patching, architecture, backups, recovery, logging, email, networking, endpoint hardening, database configuration, media handling, cloud standards, change management, etc. This one tends to be quite large and could easily be broken into several sub-domains for ease of delegation or organisation. The important part is to define, in detail, with the relevant stakeholders, how each IT activity should be done with security in mind.

5. SaaS & Business Applications: Define the state in which every internally hosted or SaaS business application should be in to ensure security, one at a time, working with the stakeholders to understand the business processes, data, and potential impacts.

6. Product and Engineering: If applicable, define all the practices that should go into your product development, product security features, hosting/engineering environment, internal and external-facing product security documentation, etc.

7. Human Factors: This is where we drive cultural change (which is related only minimally to user awareness) in conjunction with HR but also work on process engineering to reduce elements of human error in business and IT processes.

8. SecOps: This is where the stuff most people think is security goes. SOC operations, EDR, anti-phishing vulnerability management, incident notification/response, threat intelligence, forensics, etc.

9. Compliance: I do not build my security according to any 3rd party compliance standard. That is not only backwards in some ways but also likely to result in missed areas and ill-fitting implementations. Instead, my compliance is based on the application of the framework. I can then easily map my security processes and mechanisms to *any* compliance standard with minimum effort (a great business agility advantage when entering new verticals or markets). This is where those mappings happen.

10. Commercial: Here we define how we capture contract schedules that relate to security, any customer facing services like security-status portals, how we help Sales accelerate the RFI process (including any documentation to give customers), our involvement in any contractual negotiations with an impact on security, and any marketing and branding material around how our security gives us an edge. (If we think security is important, why wouldn't we have it be a brand value or selling point?)

11. Business Stubs: Links to all other departments' business processes, which should be taken into consideration when creating every component of the framework in all the domains listed above.

I typically start off building out the executive domain and getting the relevant documents approved, then create the Programme Definition layer so everyone knows how to contribute to the programme. The HR & Legal pieces (as many important other processes will depend on these) come next. Building out the rest is then mostly delegated to members of my team and collaborating other departments.

Now might be a good time to mention that, for reasons I will explain shortly, out of the dozens of activities that would fall under the IT Operations domain I always ask that the back-up and recovery piece be one of the first ones handled.

Over time, the build-out of this programme leads to close integration between departments and the involvement of those departments in security practices.

To name just a couple of examples: Access to systems can be managed automatically by HR based on role profiles and automation, systems and applications are architected and produced using templates incorporating security standards, all business processes are captured with security concerns highlighted and considered, Operational/maintenance activities like application and asset management, patching, backups, provisioning, and more become automated according to define standards, all leaving fewer gaps for potential exploitation.

The net result is greater security due to the improvement (higher quality) of the organisation's processes, not the security department's capability to firefight. This means every implemented change has a permanent effect rather than just being the latest avoided disaster.

And that is how we reverse the trend, lower the business' risk over time while needing less and less operational security resource (rather than more), save the business money, not to mention find and generate other benefits elsewhere.

I've seen everything from significant reductions in license costs, contract renegotiations, more stable applications, lower maintenance costs, higher customer satisfaction, even reduced turnover as DevOps and Site Reliability Engineers inherited fewer frustratingly hard to maintain systems. All from trying to improve the quality of processes to help, nominally, with security.

The argument of improving the business rather than "doing security" is also one I find easier to sell to executive audiences. When department heads resist, the quality argument is a powerful one. It's somewhat easy for departments to claim that it's not their job to handle security when senior leadership doesn't understand how security works. But once we've positioned security issues as being the result of poor quality, it's very difficult for those department heads to argue that they aren't responsible for the quality of their outputs.

If they insist that security isn't their responsibility, I like to ask them who they hired to lock their front door this morning.

So that, in my opinion, is the goal of a framework, and the strategic goal of the CISO.

I have yet to see anyone propose a more sustainable approach in the sense of eliminating the source of security issues in the first place.

There is however one problem with it: You must put in the work, and it takes time. Probably years.

I feel this is one reason people go for the short-term firefighting and mitigation approaches, spending most of the effort operating detection and response functions to firefight all the issues at their feet. Don't get me wrong, you need these technologies, but their greatest value isn't firefighting issues, it's giving you the end of a string to pull on to find out which business processes are responsible for the

issues. We can then rectify them so that the issues stop coming, instead of spending all our time dealing with the resulting fires.

And that's where recovery capabilities come in. I know this is a radical statement to most technical security professionals, but it allows us to, in a way, let some things burn. Or rather, risk them catching fire.

We can risk them catching fire because we know we can quickly and reliably bring them all back to life, and the time saved trying to mitigate and react to every risk or threat can be spent elsewhere *fireproofing* the business processes that build our systems, applications, and define our operations. Once we've fireproofed them, we don't need to stand guard over them all the time, which is yet more resource we can shift.

That is what risk management and security leadership should be about: using the resource we have at our disposal to deliver the highest amount of value to the organisation, and to consume as little of the organisation's resources as we can. Not to "do security" in a way that we know is ultimately ineffective (as shown by the ever-increasing frequency of breaches).

In short, we know that if the worst comes to pass, apart from a confidentiality breach, things can be put back in order. And having this capability, as you work on building out your programme to be able to stop things happening in the first place, means you can sleep at night, and focus on the future during the day.

But even from a pure technical Risk Management perspective, this idea of "letting things burn" isn't radical at all in my mind. In fact, if we include the speed of recovery in our Risk Management calculations (something that obviously *should* be done but rarely is) it affects the business impact of any particular risk.

For example, one way of decreasing the impact of *all* our risks could be not to even do anything about the threats or vulnerabilities, but merely have faster recovery.

And if we increase the scope of our risk calculations to include the root causes of the risks (business process) and start looking at a longer more strategic timeline, we find that fixing the causal business process is almost always what should be prioritised. Even more so when short term risk has been mitigated by strong recovery capabilities.

And that's precisely what we'll be looking at in our next instalment: Risk Management under the light of *business* risk, how recovery time factors in, and how it helps us shift our focus to solving the root issues rather than continually mitigating the risks they cause.

Blog (3/6) – Recovery-Based Risk Management

This might just be my most controversial instalment in this series, for security practitioners anyway.

I'm going to come straight and say it: I don't like how we do "Risk Management" in Information Security. I think that in its current guise a lot of it is of very little real value, especially from a strategic or long-term standpoint.

Let me try to clarify that somewhat. As a CISO, is it is my job to "manage" risk, or to sustainably reduce how much we have and create? I prefer the latter. Long-term, it's a lot less work for me, and a lot better for the business.

I feel we too often operate with an assumption that the business wanting to do or achieve X means Y risk, but the reality is that *how* we go about doing X greatly affects Y.

In short, it's usually possible to have the outcomes the business wants with a lot less risk (and I mean before throwing a lot of mitigating security resource at it), but it involves building the business processes with risk in mind rather than retroactively managing the resulting unnecessary or excess risk.

This is something that can be done with the proactive strategies and concepts I laid out in the Security as Quality instalment.

In other words, I don't like how the security status quo has scoped "Risk Management" and the approach that is typically used for it. Together, at a macro level, they might even contribute to us staying stuck where we are; constantly "managing" new risks rather than stopping their creation and reducing the total number of risks [to manage] at any one time, over time.

In my opinion, a lot of it stems from following two things:

1. We do not approach Risk Management from a fundamental business process angle. I.e.: We do not tend to focus on changing the processes responsible for continuously generating new risks, instead only dealing with the resulting risks without reducing the flow.

   Risk Management is often a firefighting and reactive mitigation function rather than one that makes sustainable improvements to the responsible business process in the first place.

2. The fact that even in the limited scope (as per above) of status quo Risk Management, it is risk scoring that serves as the basis on which most other risk management actions are taken. This would be fine were it not for the fact that, whether qualitative or quantitative, I don't feel we're particularly good at being accurate.

We are, as humans, often hilariously bad at determining risks and the actual causes of risks (correlation). E.g.: Pigs kill more people than sharks, cars kill more people than guns.

We also tend to simplify correlations to where A causes B, but there are often so many other factors involved that our simplified conclusion can be way off or even completely backwards. This can lead to not just poor risk calculations, but proposed actions that have little effect and can sometimes make things worse.

- Most risks are assigned arbitrary values, or even 1-5 numbers, that are disconnected from a financial business impact.

- "Quantitative" assessments, rarely are. They are still based on arbitrary assumptions that can be significantly off, whether the assessor realises it or not. I find quantitative numbers rarely consider the full complexity of the situation.

- We often lack the technical understanding of what could happen for each potential scenario. In other words, we do not know every part of every system (IPs, ports, services, operating systems, versions, installed software, patch levels, exposure windows, etc.), and all the possibilities they provide at any given time for a potential compromise.

- We rarely understand the business context and therefore actual business impact of systems. Which business process would be stopped, degraded, corrupted by a certain system being hit, or impact to another system from it, and the associated financial cost. We then need to repeat this assessment exercise for *every possible permutation* of systems being hit for any given hypothetical breach.

- There are likely many unknowns that would affect potential attack chains, leading to risk estimates being wrong by orders of magnitude. Very few are aware of every single asset of every single type they have, let alone their dependencies, interactions, roles, access profiles, what vulnerabilities they have (or could eventually have), and how that may impact lateral movement to other things.

- The state of our environment, the threat actors, and the array of exploitable situations is forever in flux. So are your risk scenarios and associated potential impacts.

- If you want to get granular, there are literally millions of potential attack vectors in even a mid-sized company. Too many to track.

- Perhaps most important of all: Good luck explaining all this stuff to a Board in under a minute!

Not surprisingly, I received strong pushback from practitioners on these views. But the fact is that in the dozens of breaches I've investigated the main cause of the breach had usually not even been identified in the Risk Register. When it was, the risk measurements, including quantitative ones, were way off.

The reality of what happened and the realised impacts (which can be counted in the aftermath of a breach) typically did not line up with what was in the risk assessment.
I'm sure many readers can anecdotally relate to what I'm saying here, especially those that have suffered a breach first-hand.

In short, I don't believe our current approach is effective and I would like to propose an alternative. Let's call it Recovery-Based Risk Management.

The basic premise is this: If you've sorted your recovery procedures properly, then your maximum recovery time is a known quantity, let's say X. The business likely knows, or can readily calculate, the loss figure for X amount of downtime.

This could be calculated for the whole business or at a department or function level, but in either case you have a maximum single incident risk impact figure.

By working backwards from the maximum negative impact instead of trying to work out every possible combination of risk, threat, and vulnerability, current and future, on systems we are not fully familiar with (if at all), and about which we don't know the full impact to business, we can work backwards from that maximum impact figure and dramatically simplify the equation.

We now have a maximum impact, which should have been set within the business' risk tolerance, due to our recovery capability. If there is an incident, the threat vector used, the vulnerabilities targeted, the sequence of the attack, how many components of the business process were affected all effectively become irrelevant. We can just assume the maximum recovery time for each specific business function and associated system(s).

We don't even need to do calculations for the impact is in business (financial) terms because we can ask the business. Afterall, it's their job to know!

Importantly, this approach is also infinitely easier to explain to management, which makes it easier to get support for. Having a single maximum risk figure is also a level of detail more suitable for executive reporting.

At this point, since we've essentially capped the possible impact of incidents, the main objective becomes reducing their frequency. Something that, as covered previously, is best done at a strategic level where the aim is to reduce the amount of vulnerable surface we generate as a business, rather than endlessly ramping up the reactive capacity to try and mitigate every attack we've made possible.

I can then apply traditional score-based practices at this broader level (rather than for individual technical risks and controls) where, due to having a definitive measurable maximum impact, they work far more effectively.

For example, rather than focusing on mitigating ("managing") my, say, 1,500 most significant vulnerabilities, I can prioritise the remediation of the handful of issues in my IT and business processes that ultimately caused them.

Are most of my issues caused by bad code? Is it my architectural practices? What about my IAM? Is it all due to lacking management support? These tend to be relatively easy to answer and prioritise compared to thousands of individual technical vulnerabilities.

It only takes a quick glance at what types of vulnerabilities we have make give a usefully accurate qualitative assessment as to how many are caused by any of these root issues. And these more fundamental and less technical issues where we can use the more traditional score-based Risk Management approaches effectively because there are far fewer variables.

They make sense when dealing with a handful of broad issues, far less so with thousands of technical risks with countless potential combinations and dependencies.

To give an extreme example, in terms of total risk reduction over time, it may even be mathematically better to completely ignore all the problems we have with our current assets and only focus on the processes that will produce new systems. Afterall, the vulnerabilities we have now will age out along with the systems they are found on, and eventually disappear. But the real point is that if we do not focus at least part of our efforts on fixing the processes that lead to those vulnerabilities existing in the first place, we will never be able to decrease how many total issues we need to "risk manage."

When Italy was the kidnapping capital or the world in the 1980's, the government did something radical to stop the problem: They made it illegal to pay the kidnappers, even freezing the assets of the victim's friends and family so that they could not raise the ransom.

When this law first passed, it was not a good time for the victims that had already been kidnapped and could not be ransomed home. But kidnappings immediately became an ineffective way to make money and all but stopped.

I am not advocating that we completely stop addressing the technical issues that we have in our environments today, but rather that some of them can be left to the safety net of recovery so that we can shift resource to where it can do more strategic good. The more we shift, the fewer risks the business will produce for us to manage, allowing ever more resource to be shifted to proactive causes, accelerating the positive trend ever further.

In simpler words, the safety net of effective recovery allows us to shift resource from chasing technical problems, to fixing the business processes that produce them. To focus on the things that will have a lasting effect on lowering the curve.

One argument I hear against this approach, where we leverage the recovery safety net to "forego" firefighting to some degree in order to focus on root causes, is that when it comes to the CIA triad, it might work with issues relating to Integrity and Availability but not Confidentiality.

This is absolutely correct.

But that doesn't mean we shouldn't do it, because this "Recovery-Based Risk Management" approach can also indirectly help drive significant improvements to the confidentiality side of things.

Firstly, resource is freed due to the simplification of the Risk Management process and the lower number of risks needing mitigation (because their *impacts* have been mitigated instead through recovery). That freed resource can be refocused on issues that can't be addressed with recovery, namely those possible breaches of confidentiality.

Secondly, a solid recovery plan requires knowing where your data is. That means that as part of the exercise of setting up your recovery capability, you tend to find out where the sensitive data is located.

This not only tells you where you might need confidentiality controls, but also where you *don't*. Systems housing data subject to confidentiality concerns make up only a fraction of most environments.

In other words, while a "Recovery-Based Risk Management" approach means treating the Confidentiality part of the CIA triad separately, it helps you decrease its scope and gives you additional resource to tackle it.

So, in summary, recovery gives us some "provocative" things to think about when it comes to how we calculate risks and prioritise resource to remediate in the most effective way, for the business, and overtime. I have yet to find a way where what I've referred to "Recovery-Based Risk Management" here couldn't at the very least augment the status quo, if not improve it dramatically and I invite everything to consider what it might mean for them. I hope you'll consider it too.

And once you appreciate that, consider the advantages and changes to risk impacts of having the fastest recovery solution in the world: Hitachi Vantara and VM2020's CyberVR. More on that in a future instalment!

Goodbye for now but please join us next time when we have a look at trends in the operational resilience space, and how rapid recovery is critical to meeting incoming regulation and the associated liability. Take care!

Blog (4/6) – Recovery & Regulation

In this instalment I'd like to look at recovery from a different perspective: How it relates to corporate liability and other impacts of present and upcoming regulations around Operational Resilience and even Privacy like GDPR.

We know from GDPR that companies often fear the associated fines more than they fear the potential breach itself. It's bad enough to have a breach and experience weeks of chaos to recover (or hopefully just hours with a top-notch recovery capability), it's a whole other thing when you're later publicly investigated and eventually fined a noticeable percentage of your earnings.

Regulation around Operational Resilience, with the potential for eventual penalties, is now starting to appear in multiple countries. This means you may now be penalised for incidents that caused business downtime, in addition to Privacy fines from GDPR, CCPA, etc.

In a nutshell you'll experience downtime (losses), have to shoulder recovery costs (more losses), and then get hit by a fine for good measure (you guessed it, even more losses). That stings.

To my knowledge, the US, Europe, UK, Australia, and multiple countries in Asia, Africa, and the Middle East are currently working on regulations around Operational Resilience. Most are, for now, focused on critical industries such as financial services and manufacturing.

Fortunately, recovery capabilities can help reduce or prevent the three areas of business and financial losses mentioned above.

Heck, it can help potentially prevent the breach from happening in the first place by empowering the Security as Quality principles mentioned in previous instalments. But in this instalment, I wanted to provide a generalised summary around the salient points of most Operational Resilience regulation, what it means to us, and how fast and effective recovery plays a role:

1. **Who is responsible for my organisation's resilience?**

   Most regulators are looking at resilience from the standpoint of the organisation and its services. It is *not* looked at from an IT-specific lens.

   It's therefore no surprise that most see the responsibility for Operational Resilience as the ultimate responsibility of the COO or equivalent.

   For example, the UK's FCA (Financial Conduct Authority) gives responsibility to what they refer to as "SMF (Senior Management Function) 24", which maps to the "Chief Operations Function". This is defined as follows:

The chief operations function is the function of having overall responsibility for managing all or substantially all the internal operations or technology of the firm or of a part of the firm.

There are however several exemptions to accountability if designated individuals have had part of the above scope delegated to them.

In most scenarios this can make the CIO accountable for the IT-relevant portions of resilience targets based on how the roles are defined to the regulator. In some cases, the CISO may also bear some responsibility, again depending on how the roles are reported to the regulator.

2. How do regulators look at Operational Resilience?

Most regulators I've seen use a model where acceptable thresholds are defined. These limits can vary based on industry vertical or even be self-defined. The idea being that organisations should define at what point an outage is causing unreasonable disruption to the parties dependent on their services as well as themselves (because that would inevitably impact the consumer of their services).

For example, it may be tolerable for a banking customer to be unable to access funds for a few hours, but a few days would be considered excessively disruptive. Some real-world examples of excessive disruption could include the relatively recent US FAA outage, the UK Royal Mail outage, and several UK and US Bank outages (some of which resulted in significant fines). Considering the impact these incidents had, it's no wonder regulators are keen to prevent incidents of that scale from happening.

Organisations must be prepared to meet the defined criteria for their various business activities. These must be reasonable self-defined values that in some cases need to be registered with the regulator. In some cases, the regulators themselves may define and impose these thresholds.

3. How does *Organisational* Resilience differ from (and complement) *IT* resilience?

When it comes to Operational Resilience, regulators are looking at *organisational* resilience and not just *IT* resilience.

While IT resilience typically supports many business processes, it's important to consider that disruption may come from non-IT failures, that the recovery of IT services may not fully restore the business process by itself, and that some business processes may not involve any IT at all.

Operational Resilience regulations' primary concern is not your IT, it is your ability to deliver services.

Conversely, switching operations to pen and paper can be perfectly acceptable from the regulator's standpoint if service levels are reasonable.

All this means that your approach to resilience, in the context of these customer-focused regulations, cannot be siloed. It must consider the full business process and all departments involved in the delivery of the service to the customer.

4. How does Cyber Insurance relate to Operational Resilience?

Cyber Insurance is a common component of an organisation's resilience strategy. It helps an organisation mitigate the financial impact of an incident by providing some coverage for expenses such as business interruption, legal fees, and data recovery costs. Its other important value proposition is that it typically gives access to expertise brought in by the insurer after an incident that can help accelerate recovery to some degree.

However, this assistance cannot be used to establish defined recovery times *before* an incident and can therefore not be relied on for defining or meeting the recovery targets looked at by regulators.

It's best to think of the insurer's post-incident services as something to mitigate unexpected situations beyond your standard recovery plan. Don't forget that there will likely also be limits in the scope of coverage and the potential risk of denied claims. In other words: Do not rely on insurance as the basis for your recovery capability, it will not tick the box.

Conversely, an organisation's existing recovery capabilities are likely to be a key factor in its ability to secure coverage, the amount of coverage the insurer will be willing to extend, and the associated premiums.

We have moved from a footing where many organisations relied on insurance to provide them with the desired resilience, to one where resilience is a prerequisite to obtaining the insurance.

In short, Insurers are more likely to cover you, for less, the better your recovery capability is.

5. What's the best way to beat resilience targets?

In a word: Speed.

Your recovery must be effective and complete. Meaning it must be well planned and consider all elements of business process (IT or otherwise) necessary to resume or continue operations and service delivery. This can entail restoring normal service or bringing up some temporary way of working, potentially with reduced but acceptable capacity.

Once you are confident that you have ways of restoring or maintaining your business functions, what matters most is simply how fast you can execute your recovery process. Operational resilience targets, as seen by regulators, are defined by the speed in which you can bring back service.

This means speed is *the* decisive factor in meeting targets to achieve and maintain compliance to these new regulations.

Preparation, planning, and testing are important pre-requisites. No complex recovery has a good chance of success without them. But they should be in place *before* the incident. Once a recovery needs to be triggered, it's all about how fast you can run your recovery process.

Now feels like a good time repeat something I mentioned in the last instalment: Hitachi Vantara's recovery solutions, when used with VM2020's CyberVR, have been tested as *the* fastest recovery solution on the market today.

That speed is going to help beat those regulatory targets, not to mention lower the financial impact to the business. Speaking for myself, it's an advantage I want.

6. What about forensics?

Quick story: My first encounter with Hitachi Vantara solutions was through VM2020's CyberVR solution. It had nothing to do with recovery at the time, I was more interested in the capabilities VM2020's Thin Digital Twins around security remediation and testing.

But I then saw that Digital Twins can have powerful uses in recovery scenarios as well.

I will discuss the possibilities Thin Digital Twins give us more fully in a future instalment, but I want to quickly mention one of them in the scope of this article: Forensics.

One of the biggest delays to recovery speed is the need for forensics after an incident. Traditionally, recovery cannot take place until forensics are completed. Something that can delay recovery by days, even weeks.

Due to their ability to replicate environments using only a fraction of the computing and storage resource, Thin Digital Twins allow for forensics to be done largely in parallel to recovery activities without the need to make a full copy of your environment (something most organisations do not have the capacity for).

This can save *significant* amounts of time and make the difference in meeting recovery targets, getting the business back up, and avoiding fines.

More on Digital Twins in a future instalment of this series.

To summarise, while we often look at recovery capability from the perspective of whether we *can* recover, the element of how quickly it can be achieved is often not as well considered. Careful planning and selection of recovery capabilities (technologies) and maximising their effectiveness through correct implementation and planning is key not just in minimising downtime, it's particularly crucial for meeting regulatory targets.

And that concludes this instalment which I hope helped frame what factors, both and in the future, matter to regulators and how to best meet them.

Join us next time as we take a closer look at what else Thin Digital Twins can do for us.

Blog (5/6) – Simulation for Remediation and Design

I personally feel like the field of storage has come a long way over the last 20 years. What used to be thought of mundane tape backups has potentially become one of the most innovative spaces in IT, with further innovation in solutions for cloud, virtual, and on-premises alike. Within it are some powerful opportunities for security practitioners I feel need to be better known and understood.

Today I want to focus particularly on how Thin Digital Twins can be leveraged to get so much more out of storage (and recovery). They were what started my deeper interest in storage from a security perspective. My reason for looking into them? To increase my ability to change the unchangeable.

Most security professionals are more than a little familiar with the frustrations of dealing with legacy systems, applications, and networks (we'll refer to all these as "systems" going forward). Systems for which they need to mitigate risks in evermore complicated ways due to their inflexibility, lack of built-in security, absence of support, and sometimes also a lack of documentation.

Note: Fairly new virtual and cloud-based systems are perfectly able to fall into the troublesome legacy category that we usually associate with more traditional on-prem systems if there was poor architecture in systems, applications, and process!

Worse still, legacy systems are frequently associated with fundamental (and critical) business processes. This often results in strong resistance by the business to us doing anything even remotely invasive with them for fear of disaster.

In an earlier instalment of this series, I talked about how a significant part of the purpose of a security programme is to define how things should be built so that new systems meet our security requirements. The same applies here, but for many of these legacy systems it's likely to be years before they are replaced. As a result, security teams spend vast amounts of resource adding compensating controls and managing risk around them.

But what if we could wipe out that hesitancy by the business? What if we could modify network configurations, code, system design, update software to versions we don't even know will work, do aggressive, potentially destructive, penetration testing, deploy patches with abandon, in a production environment context, all without needing to be concerned about the business impact, or even needing to go through change control?

Remediation that would take years could be done in weeks. Certainly in less time than it would take just to roll out another "compensating control" when not allowed to touch the systems, networks, or applications themselves.

Of course, the business impact and risk appetite of the organisation would never allow that. But what if there was no impact? What if the worst possible outcome of our intervention into these legacy systems didn't disrupt a thing?

The business wouldn't really care then, would they? But how would you do that? Those two things are diametrically exposed, mutually exclusive, you can't have it both ways.

Or can you? Enter the world of Thin Digital Twins.

In this article I am referring to Thin Digital Twins in the context of VM2020's CyberVR, but the general concept of general concept of a Digital Twin is a functional digital copy of a system (which can be a complex system of systems, such as an entire environments) that can be used for simulation, validation, and modelling… under the same exact conditions as the production one.

The Thin bit is the clever part and involves some really interesting patents. It allows you to functionally recreate complex systems (and whole environments) using only a small fraction of the resource.

And that's where the practicality comes in. Most organisations do not have the spare compute, memory, or storage to allocate in order to create full copies of their environment, but by using *Thin* Digital Twins it becomes possible to create functional replicas using only a small amount of spare capacity.

This means environmental changes can be tested under conditions equivalent to full production by using clones of production that replicate all environmental factors.

What does this mean to us?

Well, in the case of legacy systems, we can now generate a materially identical instance of any system (which we can pull from storage as our backups are essentially "dead" copies of our environments waiting to be brought to life) and do anything we want to them without consequence.

We can try integrating new functionality, code changes, testing the impacts of updates or patches, or even rip out a system entirely and replace it with something else to see if there's any impact to the business process. We can alter one system at a time to see if there are impacts with other replicated systems, or we can make comprehensive changes to our whole environment.

We can find out if configuration hardening changes have any negative impact on the system's needed functionality, or even aggressively pen-test "production" systems without any fear of consequence and gain insights we never could before.

In previous instalments we talked about how reactive security practices (and having to put in "compensating controls" is definitely one of those) keeps us from working on the long-term strategic changes. The kind that can sustainably improve an organisation's security posture to not only reduce risk but also reduce the cost to the business and our workload.

We discussed how outputs from business processes resulted in more and more risk to manage (and, consequently, work) due to defects caused by a lack of security thinking and integration in the creation of those processes.

If you think of removing the security defects of these processes like picking up stones on the road that are damaging passing cars, then legacy system are a bit like boulders; We can clean up the smaller things relatively easily and make sure new stones don't end up on the roadway, but the boulders can't be moved and we're going to have to do all sorts of things to keep them from causing incidents.

Thin Digital Twins can give us the leverage and power to clear these boulders, often our biggest obstacles, much faster.

They're also powerful tools for any kind of modelling of new systems, or any changes or interventions to existing systems that usually incur some resistance.

They help us prevent more from happening by enabling us to reduce existing security risks more aggressively (without us causing business risks), and better ensure that new systems aren't introducing new ones.

In a breach recovery scenario, the ability of Thin Digital Twins to replicate an environment with fractional resource allows us to perform forensics *in parallel* to recovering the full environment.

This means we can immediately proceed to recovery and then only focus on cleaning up the parts of the environment that need to be. In the past, forensics and clean-up had to happen before the recovery could run because there simply wouldn't be enough spare computing resource to run a full parallel environment.

This parallelisation of forensics and recovery saves time, and as we know, during a recovery operation, time is everything.

There's another element of parallelisation that is unequalled with the VM2020 offering. It relates to just how fast data can be pulled out of your backups.

In past instalments I've mentioned recovery times of four hours to give an example. These are not what most people would consider realistic figures due to the incredible I/O loads involved in moving data out of immutable storage. Recovery times of two to three weeks are likely more typical, though they would include several days of forensics before recovery (restoring data) was started in earnest.

So why have I been making arguments with unrealistic figures?

Well, they *aren't* unrealistic anymore, because VM2020's solution is so effective at optimising this process that in recent tests using Hitachi Vantara's Ops Center Protector the recovery of more than 1,500 virtual machines with over 100TB of data was achieved in 70 minutes.

70 minutes. Think about how being able to recover that much in such a short time changes our risk management calculations! Think about how much that limits the impacts of any risk, of any disruption becoming a major one.

Note: While those figures were achieved on virtual machines, similar performance is possible both for cloud and physical systems due to patented technology allowing CyberVR to instantly virtualise and "devirtualise" systems and apply the same process.

As mentioned previously, that makes the combination of Hitachi Vantara and VM2020 the fastest recovery solution in the world. Read for yourself using the links below the article.

https://www.hitachivantara.com/en-us/pdf/solution-profile/worlds-fastest-ransomware-recovery-from-immutable-snapshots-vmware-environments.pdf

https://www.hitachivantara.com/en-us/insights/using-thin-digital-twins-to-gang-up-on-ransomware.html

And that seems like a pretty good place for me to end this instalment. I'll leave you with a few links to find out more about VM2020 and hope you'll join me for our final and perhaps more conventional instalment where we have a look at what a good recovery looks like and how we should prepare to be able to execute it and leverage all the benefits at our disposal.

Blog (6/6) – The Importance of Recovering Well

Welcome to the final instalment of this series on how recovery can enable us to progress how we do information security.

We've seen how we can get proactive, shift focus and resource towards that approach thanks to the safety net of recovery, use Digital Twins to tackle some of our toughest legacy challenges and to accelerate forensics, what the regulatory considerations are, how speed *really* matters, and much more.

Most of these capabilities have one particular caveat though: Your recovery capability needs to be *good*. It must work well at recovering your business processes in the real world. All of them. You cannot work towards implementing a proactive programme tomorrow if you aren't confident in your ability to recover today. The distraction will consume you.

And this is where we address the more traditional or conventional aspects of backups and recovery. You may recall that one of the first things I said in this series was that I was not an expert in storage and recovery. Fortunately, there's plenty of best practices out there for us to work from.

And this being 2023, what blog wouldn't be complete without us asking ChatGPT a few things? Let's have a little fun with our favourite artificial friend and get their thoughts on recovery:

Organisations rely heavily on digital data and systems to run their operations. Losing this data due to system failure, cyberattacks, or natural disasters can have a devastating impact on the organisation's ability to function. Therefore, it is essential for organisations to have an effective backup and recovery plan in place.

Well thanks for that, Captain Obvious. So how do we achieve that? What's the process?

1. *Identify critical data and systems.*
The first step in planning for effective backups and recovery is to identify the critical data and systems. Not all data and systems are equally important to the organisation. Therefore, it is essential to identify the critical data and systems that need to be backed up and prioritised for recovery in case of a disaster.

Organisations can use the following criteria to identify critical data and systems:

- *Importance to business operations: Identify the data and systems that are essential for business operations. This could include customer data, financial data, inventory data, and other data that are crucial to the organisation's day-to-day operations.*

- *Regulatory requirements: Identify the data and systems that are required by law or regulation. For example, the organisation may be required to keep certain financial data for a specific period of time.*
- *Recovery time objective (RTO): Identify the data and systems that have a short RTO. RTO is the maximum amount of time that an organisation can afford to be without a system or data before it starts to impact business operations.*
- *Recovery point objective (RPO): Identify the data and systems that have a short RPO. RPO is the maximum amount of data loss that an organisation can afford in case of a disaster.*

Ok, that sounds pretty decent, applying long-established common-sense principles. Though it's important to note that just because the IT supporting a process is down doesn't mean it can't be done manually. And it's not just how quickly something starts having an impact, but how big. The $10,000,000 impact at midnight due to a failed batch job is a bigger worry to me than the $10 per minute we're losing on some live activity, for example! So do consider those things in your prioritisation. Sounds basic but I've seen this not done properly in a number of cases.

Another thing I would like to point out is how ChatGPT keeps bundling together "Data and Systems."

When it comes to backups, I would personally advise focusing on the data first as there may not be a one-to-one correlation between data and systems, then the business process(es) they're used in, then the systems supporting those processes.

When it comes to *finding* the data however, looking at the business process first is likely your best option. Experience tells me this is the best way because the business process dictates all.

That doesn't mean it may not be collecting excess and unnecessary data. But that's a separate architectural (and possibly liability) issue beyond the scope of our discussion here.

Ok, next step?

2. *Determine backup frequency and retention period.*
Once the critical data and systems have been identified, the organisation needs to determine the backup frequency and retention period. Backup frequency refers to how often the data and systems are backed up, while retention period refers to how long the backup data is kept.

The backup frequency and retention period should be determined based on the criticality of the data and systems, RTO, and RPO. For example, if the critical data and systems have a short RTO and RPO, they should be backed up more frequently and kept for a longer retention period.

Organisations can use the following guidelines for determining backup frequency and retention period:

- *Daily backup: For data and systems that are critical to business operations and have a short RTO and RPO, daily backups should be performed.*

- *Weekly backup: For data and systems that are less critical, weekly backups may be sufficient.*
- *Monthly backup: For data and systems that are not critical, monthly backups may be sufficient.*

The retention period should be determined based on regulatory requirements and business needs. For example, if the organisation is required to keep financial data for seven years, the retention period for financial data backups should be at least seven years.

More common sense here. But again, we see the generalised mention of "data and systems" which should probably be replaced by a consideration of the business process that needs the data. In fact, ChatGPT seems to have a habit of glancing over the business putting other things first. This is a symptom of the status quo and something we need to change if we are to understand our organisations well enough to protect them.

One consideration ChatGPT hasn't mentioned is the importance of the timeliness of a specific type or subset of data. Names and birthdays may be static, but if your organisation's role is to track trends then historical data, with a full record or high enough sampling rate may also be important. This touches on the subject of full backups versus incremental ones, and the possible use of transaction to not just bring data up to date but have its full history.

Something that may also be worth pointing out is that in some cases we may need copies of data significantly fresher than one day old which is the most frequent suggested here. In fact, even minutes could be too much. Ensure you will have the data you need for each purpose.

All could require different approaches and planning to ensure we have a copy of the relevant data to support the recovery strategy for each business process.

Now we need to pick a solution (or solutions) that meets our needs and start working out and documenting the backup [and recovery] processes accordingly.

3. *Implement backup and recovery policies and procedures.*

Backup and recovery policies and procedures should include the following:

- *Backup frequency and retention period: Document the backup frequency and retention period for critical data and systems.*
- *Backup and recovery methods: Document the backup and recovery methods used, including tape backup, disk backup, cloud backup, or hybrid backup.*
- *Disaster recovery plan: Document the disaster recovery plan and the steps to be taken in case of a disaster.*
- *Roles and responsibilities: Document the roles and responsibilities of the backup and recovery team and other stakeholders.*

Backup and recovery policies and procedures ensure that the backup and recovery plan is followed consistently and helps to reduce the risk of data loss or system downtime.

After all that we should now have a documented backup and recovery process.

But how effective is it?

Next up we need to test if what we've thought up and documented will *actually work* in terms of the recovery of business processes. You'd be staggered at how few organisations have done this well or at all only to get caught out in a big way.

It's critical to define also not just how we will recover systems and data, but how the people performing the business processes will resume using them. There's nothing quite like having worked all night to get systems ready again for the business day, only to have the entire workforce twiddling their thumbs in the morning because no one knows how to start or connect to the business application as no one's ever had to do it before in living memory, it was always just "on".

4. Test backups and recovery
Once the backup and recovery solutions have been implemented, it is crucial to test them regularly. Testing ensures that the backups are valid and can be used for recovery in case of a disaster. Testing also helps to identify any issues or gaps in the backup and recovery plan.

Organisations can test backups and recovery in the following ways:

- *Partial restore: Restore a portion of the data to ensure that the backups are valid and can be restored.*
- *Full restore: Restore all the data to ensure that the backups are valid and can be restored to the original system.*
- *Simulated disaster: Simulate a disaster to test the recovery process and identify any issues or gaps in the backup and recovery plan.*
- *Tabletop exercise: Conduct a tabletop exercise to test the backup and recovery plan and identify any issues or gaps in the plan.*

Here I'd add to not just test the recovery solution but also the plan, including the business processes and people running them to make sure things align nicely to the business and not just the IT.

Identify any issues, adjust your plan and documentation, and retest until everything goes smoothly.

There's a slight issue that the status quo best practices don't always mention; testing this stuff is *really hard*. It's hard because it's traditionally extraordinarily time consuming, disruptive, and limited by spare resource, both in terms of manpower and computing resource.

Doing testing requires significant computing and storage resource the business doesn't likely have spare, which can result in partial and fragmented testing which may let us down in a real-world scenario. Having to do it carefully due to the associated risks also exacerbates the human effort required.

But it's another great use case for Thin Digital Twins like those leveraged by VM2020's CyberVR. Full scale recovery can be tested with far less worry. Meaning you can not only make sure you test everything, but you can test it faster and quickly reset test cycles and repeat until your recovery process is successfully proven. (Not to mention perform your forensics in parallel in the case of a breach.)

Remember that all these principles apply to all data and systems, whether physical or virtual, on-premises or cloud. While people usually assume Digital Twinning solutions only work on virtual machines, I'm glad to say CyberVR gives us options for all of these scenarios.

Remember when we said that Hitachi Vantara's Protector, when used in conjunction with VM2020's CyberVR could recover 1,500VMs with a petabyte of data in 70 minutes?

Well, I thought I'd ask ChatGPT how long it thought that would take. "Days" was the answer I got back. And once I added the elements of immutability, validation, forensics, and different elements of cloud, virtual, and physical on premise, it became "Several days or weeks."

You can see why this is exciting and a big deal when it comes to recovery. And, as a CISO, how it enables me to deliver a strategic and proactive security programme to improve my organisation's *inherent resilience* (where we don't get knocked down in the first place), thanks to the safety net and the change in priority it allows due to *Recovery-Focused Risk Management.*

Furthermore, from a from a security perspective, it's essential for your backups to be not only encrypted, but also immutably stored so that they cannot be compromised either from a confidentiality or integrity standpoint.

Another important element, that also significant impacts our ability to do Recovery-Focused Risk Management with quantitative accuracy, is how consistently we can recover from incidents (and we can even test this thanks to Thin Digital Twins).

This accuracy means we can provide more assurance to our Board and make us exponentially more likely to meet our recovery targets. It helps us not just correctly prioritise (budget) business resources for risk management according to the financial *business risk*, but also improves our ability to set and meet regulatory targets should the worst happen.

One thing I want to highlight is that there is a difference between traditional backups and the kind of recovery we have been discussing in this blog, as these differences may not be clear for security practitioners who don't live and breathe storage.

Traditional backups are great for recovering files and data, but don't provide a capability around recovering systems and business processes. In other words, they do not focus on recovering business services, which has traditionally required a lot of additional effort.

That doesn't mean traditional backups are "bad". To the contrary, sometimes all you want to do is back-up and restore files, and it is something those methods are time-tested and exceedingly good at, just keep in mind the difference between recovering critical business services and merely recovering or restoring files.

The chart below can serve as a guide which highlights and explains the differences:

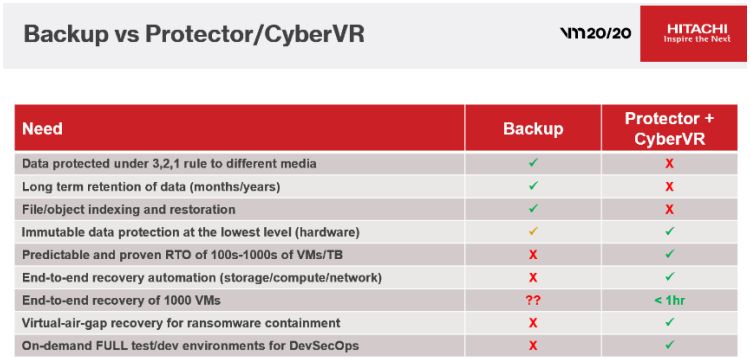

Backup vs Protector/CyberVR

VM20/20

HITACHI
Inspire the Next

Need	Backup	Protector + CyberVR
Data protected under 3,2,1 rule to different media	✓	X
Long term retention of data (months/years)	✓	X
File/object indexing and restoration	✓	X
Immutable data protection at the lowest level (hardware)	✓	✓
Predictable and proven RTO of 100s-1000s of VMs/TB	X	✓
End-to-end recovery automation (storage/compute/network)	X	✓
End-to-end recovery of 1000 VMs	??	< 1hr
Virtual-air-gap recovery for ransomware containment	X	✓
On-demand FULL test/dev environments for DevSecOps	X	✓

Some parting words to my fellow security practitioners in this, the final instalment, of this blog:

Like me, most of us are not experts in storage and recovery. But we must accept and appreciate that it is a highly complex discipline, likely rivalling that of security.

This instalment only covers the utmost basics, and I would recommend leveraging expertise from specialised consultants (such as those at Hitachi Vantara), due to the experience and expertise needed at a business and IT operations level as well as in terms of the degree of product and technology knowledge needed to get it right.

But whether we use internal or external expertise to get it right, once it is we gain immense possibilities in changing how we do Risk Management, and how we can approach the work of securing and reducing the risk to our organisations. Moving away from mitigating and firefighting technical issues caused by business processes (including IT), to instead affect those processes themselves, building security *into* them rather than adding it to them wherever possible, and creating sustainable improvement in our organisation's security posture similar to the trend we saw in the aviation sector in this blog series' introduction.

As technology becomes ever more important in everyone's lives, we have an opportunity to make a *lasting* mark, one that truly matters.

Thank you for reading.