Synechron

Cyber Security and the growth of untrusted infrastructure and hybrid workforces

Q&A with our SMEs on: Cloud Native Security

Today's mass usage of cloud services presents an enormous challenge for security teams from highly regulated industries. As financial enterprises shift their focus to DevSecOps, security capabilities and skillsets are now expected to be embedded as part of the DNA of every application and platform. This is triggering a cultural shift in the adoption of cloud native security practices in financial services firms.

We spoke to Synechron's and Synechron's SMEs about cloud native security and what it means in a post-pandemic world where the workforce operates in the 'new normal' whilst malicious parties take advantage of the vulnerabilities.

Eklove Mohan

Sr. Director-Technology
Ashburn, Virginia, USA

As enterprises pivot to building business critical systems on untrusted public cloud infrastructure, how will cyber risk be balanced with the simultaneous shift to a hybrid workforce on untrusted devices?

'Untrusted' is the perception but the reality is 'unawareness'. Look at any of the cybercrimes and you will see that the cause of the breach is mostly due to 'misconfiguration'. Cloud has a 'shared responsibility' security model. The Cloud Service Provider's (CSP) responsibility is to provide infrastructure and ensure security 'of' the cloud but not 'in' the cloud, i.e., the application and data have to be secured by the customer with either the tools provided by the cloud provider or a third party.

Some of the key factors that should be considered to balance cyber risk include:

- **Utilize 'least privilege'** – By default, all access to cloud resources is denied. You only enable a resource if it is required by an application.
- **Use certified images** – Restrict the use of public images. Ensure that the security team within the organization tightens (blocks ports, installs certified versions of tools/applications, disables telnet, etc.) and certifies the OS images used by the application. This will ensure consistency, and if a flaw is identified, it gets fixed for all.
- **Select cloud native solutions over third-party tools** – Most organizations prefer to keep things cloud agnostic. Hence, opt for third-party tools. However, if the organizations have pretty much decided on a cloud provider, it is preferred that they go with their cloud native services. We have seen that in terms of security, the implementation/patches by CSPs are most up to date as compared to third-party tools.
- **Choose virtual desktop services** – One of the biggest headaches for any organization is protecting their own assets from misuse. With physical devices, like laptops/desktops, there is always a chance of loss or theft. A malicious user many not easily be able to log in to the stolen device, but they can surely read the data from the hard disk easily. With virtual desktops, the company's data is not sitting on the machine. Also, virtual desktops are connected to the organization's active directory, hence the security is centralized and controlled by an elite set of security professionals within the organization.

Are static and dynamic code scanning tools enough to assure security principles in the Software Development Life Cycle (SDLC)?

When it comes to security, nothing is ever enough. Static and dynamic code scanning tools will just help to ensure that you have the basic scenarios covered and the hacker has to sweat a little to get into the application/system. In the recent past, organizations have understood the fact that they may never be able to totally protect their application or infrastructure from being hacked. Therefore, what they are focusing on is how quickly can the attempt to hack be identified and then remediated or blocked. Due to this, we have seen a series of tools being used during the SDLC phase which includes: Runtime Application Self Protection (RASP), User and Entity Behavior Analytics (UEBA), Web Application Firewall (WAF) shielding, obfuscation, etc. Consider it like an onion which has all these layers of security and at the core is the application. To get to the core, the hacker needs to peel all these layers. This Defense-In-Depth (DND) technique ensures that either the hacker gives up (unlikely) or the organization catches them before they can get to the core and takes preventive actions before the damage is done.

Abhilash Panickar

Sr. Director-Technology
Ashburn, Virginia, USA

We've been evangelizing DevOps for several years now. Why don't we see an order of magnitude impact on value delivery? Are Agile DevOps teams bottlenecked by production change gates and risk assurance? Where is the chain of trust failing and why?

Software release management practices in most large enterprises continue to rely on the traditional multi-layered and often manual change management processes. One would expect that as a result, outcomes become predictable, risks are reduced, release cycles are consistent, etc. However, the reality is that most of these big bang releases typically result in extended downtimes, system outages, support calls and eventually a scramble to roll things back to bring systems online before the start of business hours. It might appear that the lack of DevOps adoption is to blame for this predicament, but it is 2021 and many enterprises have already jumped on the DevOps bandwagon due to its overwhelming popularity as an industry trend. The issue, however, seems to be a result of the lack of maturity in the DevOps processes that does not provide the required reliability and predictability of assurances that production environment custodians and operators demand as part of their day-to-day operations. This has resulted in a hybrid approach in most organizations where deployment pipelines and high automation is applied for releases in the lower environments, with the traditional release process being followed for the higher, more sensitive environments.

How do we go about mitigating some of these issues and challenges? Organizations need to move past the experimentation and isolated targeted adoption phases into a more comprehensive approach towards adopting the DevOps model and strengthening it with the principles behind the 'Continuous Delivery' approach to bring more reliability and trust into the automated, low touch release processes. It is a proven model and many global organizations have successfully implemented it at scale. Most of the technology required to enable such capabilities are now available either in open source, or as native cloud platform capabilities, or as SaaS solutions.

Some of the key principles to adopt include:

- DevOps culture which looks at automating every step of the release process
- Automated integration and smoke tests
- Many small incremental releases versus one big-bang release
- Deployment architecture and release processes that limit the 'Blast Radius'
    - Dial-Up Weighted Routing – Release new versions to a smaller userbase and then dial-up/rollback as needed - 5% > 10% > 25% > 50% > 75% > 100% users
    - A/B Testing
- Automated metrics-driven release gates that control the release rollouts and rollbacks based on actual traffic data
- Blue/Green, Red/Black for instant cut-over and rollbacks
- Synthetic transactions to uncover potential issues automatically coupled with self-healing/recovery capabilities
- Smart operational monitoring with anomaly detection using ML

How are one-time security & compliance assurance gates failing the enterprise in both agility and risk assurance? And why is now the time to prioritize fixing their shortcomings?

As organizations become more agile with ever shorter release cycles (weeks/days versus monthly/quarterly), the security and compliance model followed traditionally -- with review checkpoints prior to the release -- is mostly outdated and ineffective. The Security and Compliance domain has traditionally been looked upon as a specialized skillset managed centrally by an independent team; vertically focused, rather than being decentralized across the different teams in the enterprise. To be effective in the rapidly changing IT landscape which has embraced the API-First, Cloud-First, Automate-Everything mentality, Security and Compliance processes within the organization need to adapt, as well as to provide a more continuous form of evaluation and verification instead of a one-off review before every release.

What are some of the changes that enterprises can adopt to become more effective?

- Embrace Security First along with API and Cloud First -- Security and compliance are not an after-thought. Instead, teams should embrace the shift-left mentality and make it an integral part of the team's responsibility – 'design it, build it, run it'. The teams can be supported by the horizontal Security COE team to provide the necessary guidelines and guardrails for them to be successful.
- Implement automation – Automate compliance testing on a continuous basis by embedding security and compliance testing tools in the build and deployment pipelines. DevSecOps is the industry term for this unified approach that brings together operations, security and app dev in a collaborative framework.
- Augment with manual reviews – Not every threat vector can be analyzed and verified using the automation tools. Use a periodic manual review process to strengthen the tools-driven compliance process.
- Collect and analyze everything – Implement a continuous monitoring program driven by real data to provide the required information necessary to support decision making and compliance assessment.
- Enable security trainings – Security is a continuously evolving landscape requiring periodic upskilling/reskilling to make the teams aware of the emerging and on-going threats.

Chris
Zanelli

Associate Partner
New York, USA

How to determine if your enterprise is prepared for the inevitability of cyber exploits? Why have your investments in high availability and fault tolerance not prepared you for recovery from Cyber scenarios and exploits?

Cyber threat and exploit scenarios differ from traditional disasters in several key ways, often requiring additional recovery capabilities not handled by application failover or traditional disaster recovery site failover strategies. The key challenge to recovering from cyber exploits lies in the fact that Cyber scenarios have complex, atypical patterns that present themselves with high variability across:

1. the time of the exploit
2. the time the impact is realized
3. the time of detection

Given that each of these critical timeframes can vary significantly, it poses a significant challenge to answering the two fundamental questions to structure recovery:

- When was the last known good state?
- How do I recover to that last known good state?

Executing a point in time recovery against the moving parts of infrastructure, data sources, data schemas, application code and configurations can be quite complex and goes far beyond how enterprises organize responses to disaster scenarios and structure their change management capabilities.

How should enterprises protect business critical and sensitive data? Must enterprises maintain air-gapped backups that are protected from active compromise? How can an enterprise prepare itself against latent attacks and protect data integrity?

Cyber Security and Recovery from cyber exploits requires a high level of confidence that the desired recovery state is 'known to be good' from both an operational perspective and from a data integrity perspective. As traditional resiliency and disaster recovery has focused on minimizing downtime and outages through active replication and geographically diverse 'active/active' sites, what we have learned about cyber exploits is that they aggressively seek to move laterally across infrastructure and data.

Highly available systems with near-real-time replication increase the likelihood of malware, ransomware, or other exploit code replicating itself across your distributed environments. Cyber exploits tend to exhibit latency in two critical timeframes which impede identification and response:

- Latency between the point of impact and remediation, which represents your actual customer/business exposure
- Latency between the time-of-exploit and the time-of-impact, which represents a period of unknown exposure

As the drift between the time-of-exploit and the time-of-recovery increases, not only does the potential for data loss increase but the likelihood of successful recovery also diminishes greatly. Successful recovery of data assets becomes a function of how long you retain air-gapped backups for your critical data sets.

As an enterprise's data infrastructure is a massive target for bad actors, how can you ensure that you know where your all your golden source data is and how to recover it?

As large enterprises seek to enhance their capabilities around air-gapped backups, data retention, and point-in-time recovery it is imperative to ensure that the investments that are being made and the priorities placed around data recovery are centered first and foremost around golden source and authoritative data sets. A well thought out Cyber Resiliency program must spend a significant amount of time making sure that the most critical and hard-to-replace data sets are protected first, whilst data sets that can be rebuilt or derived from golden source copies can be a secondary priority.

The benefits of establishing this level of data governance and lineage will be significant not only for enhancing recovery from cyber exploits, but will also be significant in establishing enterprise data management capabilities across critical data sets and the lineage of data re-use across the organization. Organizations looking to reduce duplication of data assets, eliminate the costs of data duplication, enable better data warehousing, reduce their compliance footprint, and improve overall data quality will also find significant improvements in their Cyber Security posture.

How can enterprises start to understand the impacts of Cyber threat scenarios on their lines of business? Is now the time to enhance a BCP strategy for Cyber Resiliency if it's centered on executing site failover patterns only?

As mentioned previously, traditional resiliency and disaster recovery efforts have focused on minimizing downtime and outages through active replication and failover to geographically diverse 'active/active' sites. While ensuring a rapid recovery to localized failures, this strategy can be counter-productive in a cyber exploit.

Today's exposures and cyber scenarios require the enterprise to invest both in High Availability and full stack, point-in-time recovery solutions. The capabilities employed for Disaster Recovery need to ensure there is an adequate data retention period and full stack recovery plan on hand for point-in-time recovery that will pre-date the initial exploit. Businesses that fail to invest in this area will be highly exposed to ransomware and extended outages caused by the inability to identify and restore to a known good state, and quickly resume services in the face of an outage or an ongoing public exploit.

Incorporating Cyber scenarios in your Disaster Recovery planning, testing, education, and other exercises is crucial to understanding your ongoing capability gaps and execution maturity, and ultimately reinforcing what the mechanics of cyber recovery execution require from both tooling and expertise perspectives.

About our SMEs:

Eklove Mohan
Sr. Director-Technology
Ashburn, Virginia, USA

Eklove plays a variety of roles in Innovation, Research & Development, Cloud Initiatives, DevOps, Technology Evaluation, mentoring and training of the young IT generation. During his time at Synechron, he has had extensive exposure and experience working with cutting-edge technology through the delivery of tactical and strategic solutions to clients. He also has had the opportunity to work with some of the most inspirational thought leaders in the industry.

To know more about Eklove

https://www.synechron.com/profile/eklove-mohan

Abhilash Panickar
Sr. Director-Technology
Ashburn, Virginia, USA

As a software architect, Abhilash spends most of his time analyzing business requirements, articulating solutions, driving the decision-making process for key stakeholders and working closely with development teams to translate design and vision into reality.
He has developed expertise in implementing proprietary solutions and building apps using various BPM and SOA platforms in On-Premises, Cloud and Hybrid deployment models using different development methodologies as well as different architecture & governance frameworks.

To know more about Abhilash

https://www.synechron.com/profile/abhilash-panickar

Chris Zanelli
Associate Partner
New York, USA

Chris is a versatile and well-rounded technologies expert with 15+ years' experience in the Financial Technology Services industry with expertise in both commercial technology offerings and building end-to-end enterprise technology platforms. He has a proven track record of high impact delivery that emphasizes business outcomes through product ownership, agile ways of working, data management, and software development automation. Coupled with an ORIE background with a focus in linear and non-linear programming and Neural Networks, he brings practical experience in scaled delivery to the emerging applications of Artificial Intelligence and Machine Learning.

To know more about Chris

https://www.linkedin.com/in/czanelli/

www.synechron.com