EBOOK

rubrik

Navigating Through the Clouds:

Effective Data Security in Multi-Cloud Environments

Table of Contents

The State of Cloud Cyber Resilience

Data growth favors Cloud & SaaS. According to our own Rubrik Zero Labs report, new data creation is predominantly in the public cloud and SaaS, which makes sense if you think about where most experimentation and innovation related to data is taking place today. Where in Rubrik observed organizations, on-premises data grew 20% in 2023, Public Cloud data grew by 73%, and SaaS by 145%

Malicious actors have also understood this and are targeting this extended attack surface—they tend to dig where they are most likely to strike gold.

Using the Public Cloud can introduce security blind spots for traditional security tooling as the workload and data types don't necessarily mirror your on-premises estate. For example, 70% of all data in a typical cloud instance is object storage, like Amazon S3, or Azure Blob, and data within those buckets tends to contain a lot of sensitive information. The result is that Cloud and SaaS are becoming fertile ground for malicious actors. Public Cloud (66%) and SaaS (67%) are the leading attack surfaces today.[1]

For you as a defender, it all starts with visibility across your assets, regardless of where they happen to be located. You cannot secure what you cannot see, especially in the multi-cloud visibility, which is hampered by a myriad of heterogeneous systems.

According to industry analyst Gartner[2], we'll spend north of 200 billion dollars globally on cybersecurity worldwide this year, no doubt preventing a massive amount of cyber incidents. But simultaneously, we continue to see successful breaches reported daily. It is estimated that a successful ransomware attack takes place every 40 seconds[3].

The reality is that nothing is 100% secure, and cyber-attacks will likely keep happening. People will click on something they weren't supposed to, no software is free of bugs, some of those bugs expose security flaws, and some of those flaws can be exploited. Even the most well-protected companies can fall victim to zero-day attacks.

1 https://www.rubrik.com/content/dam/rubrik/en/resources/report-review/rpt-zero-labs-4.pdf
2 https://www.gartner.com/en/newsroom/press-releases/2023-09-28-gartner-forecasts-global-security-and-risk-management-spending-to-grow-14-percent-in-2024
3 https://www.forbes.com/sites/forbesbusinesscouncil/2023/04/06/understanding-ransomware-attacks-and-how-data-centers-can-protect-themselves/#:~:text=According%20to%20Cobalt%2C%20ransomware%20is,quantify%2C%20as%20many%20go%20unreported

“Everything fails, all the time” is a famous quote from Amazon’s Chief Technology Officer Werner Vogels[4]. I would add “Everything is attacked, all the time”.

Especially when changes are introduced in an environment, those changes tend to lead to small gaps attackers can find and exploit.

“Everything fails, all the time” – Werner Vogels, CTO Amazon

We used to say it is not a matter of IF, but WHEN before a cyber incident happens to you. Today we see that it is not only a matter of IF, but rather HOW MANY TIMES you may stand to fall victim.

When we approach cyber security with an “assume breach” mentality, it can remove mental roadblocks and help shift our approach to embracing both prevention and recovery to ultimately achieve cyber resilience. How quickly can you bounce back?

4 https://thenextweb.com/news/werner-vogels-everything-fails-all-the-time

The Cloud Attack Surface

While cloud technology is not new, and many organizations have been on their cloud journey for years, hyperscalers continue to evolve with new features and services. This rapid evolution and growth makes it difficult for organizations to keep up, inadvertently introducing security vulnerabilities and expanding your attack surface.

Your attack surface is the sum of vulnerabilities, pathways, or methods, sometimes called attack vectors, that malicious actors can use to gain unauthorized access to your cloud environments, including a growing amount of sensitive data, to carry out a cyberattack.

As organizations increasingly adopt cloud services, including multi-cloud and SaaS, the associated attack surface is becoming larger and more complex.

According to Proofpoint[5], almost all cloud tenants were targeted every month in 2022, and 2 out of 3 were successfully compromised.

Even state-sponsored adversaries are changing tactics according[6] to the UK's National Cyber Security Centre and the Five Eyes partners. In a recent joint advisory, they pointed out that malicious cyber actors linked to Russia's Foreign Intelligence Service (SVR) are adapting their techniques in response to the increasing shift to cloud-based infrastructure.

A benefit for malicious actors in cloud environments is that with the right credentials, access is easily achieved.

Today, hackers will attempt to log into your environment more than try to break in. Valid credentials are readily available, either accidentally left on sites like GitHub or Pastebin, or cheaply purchased as part of identity acquisition campaigns on Telegram or Discord, often including valid session cookies to overcome additional security layers like Multi-Factor Authentication (MFA).

5 https://www.proofpoint.com/us/resources/threat-reports/human-factor
6 https://www.ncsc.gov.uk/news/uk-allies-expose-evolving-tactics-of-russian-cyber-actors

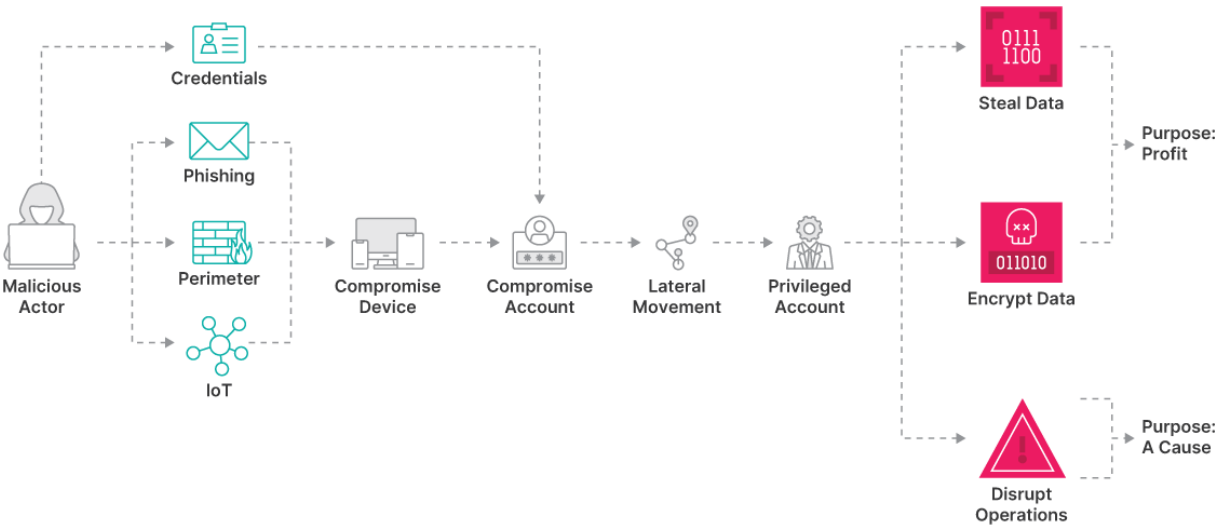

Consider the diagram above, a malicious actor can leverage valid credentials of a compromised account to log into your cloud environment, maintain persistence in the environment move closer to valuable assets, move from one stolen identity to the next, and look for the right one to execute their desired outcome. No need to phish an individual and compromise their device, no need to bypass a firewall, or find an exploit through other means to get in. The additional benefit is that traditional security tooling will be less suspicious about valid credentials and any activity they generate.

In Public Cloud, it becomes a lot harder to build thicker walls and broader moats around your apps and data. Systems are highly distributed, and native services store and move data in all kinds of interesting and novel ways. Additionally, the default security controls you might have in one Public Cloud don't necessarily translate to the other.

When it comes to data, the hyperscalers have always worked on the basis of a shared responsibility model, whereby the ultimate responsibility for data, data backup, and data security remains with the end-user. So when adopting a cloud or multi-cloud approach, we need to consider whether we have the capabilities to recover from a ransomware attack. Do we even know the sensitive data and associated risks we have across those cloud providers? In case of exfiltration, can we report on this in a timely manner to the requesting authorities? How realistic is preventing data theft, and do we understand who is ultimately responsible for data recovery across the multi-cloud, especially in light of increased regulation?

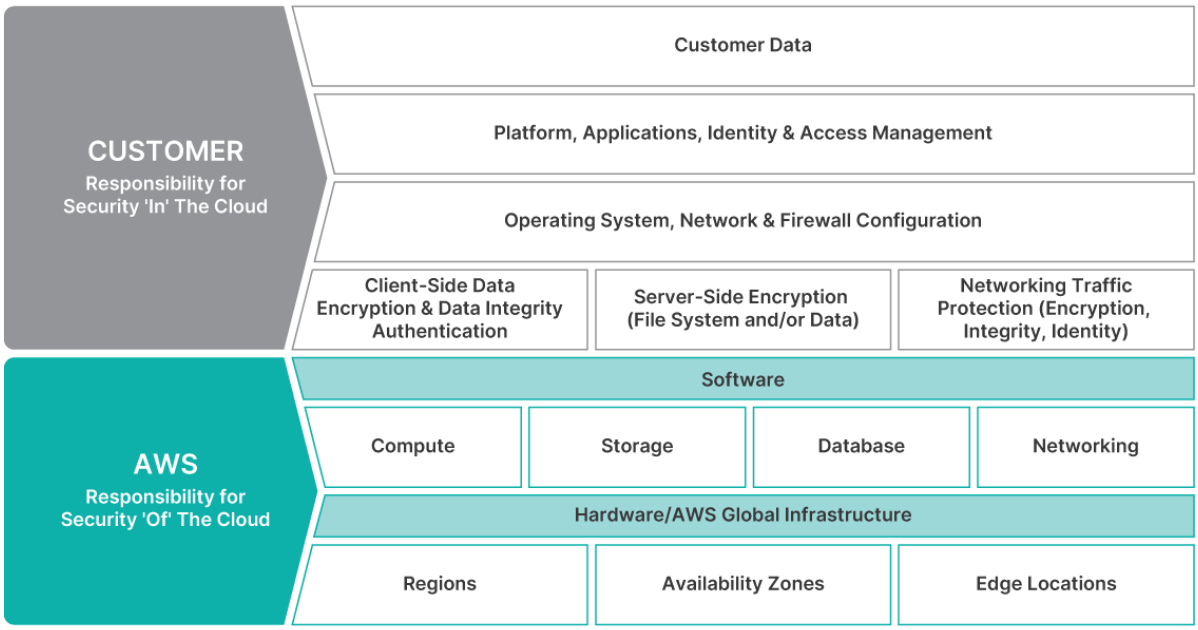

More Clouds, More Problems: Are You Resilient with Cloud Native Tools?

Not only do you have multiple clouds, but you also have a wide range of services across IaaS, SaaS, and others. Some of these assets contain sensitive and regulated data.

Even in the same cloud, not all services are created equally. Think about providing BCDR for all your cloud services in a single cloud. It will rapidly start to resemble BCDR for on-prem environments, where you have to contend with various applications and approaches.

Your teams working in this multi-cloud environment have specific needs and tools. The pressure on those teams also comes from multiple angles. They operate miscellaneous tools, oftentimes not fit for purpose from a cloud perspective, and have to do this across many environments. The external pressure being leveraged by malicious actors are potentially unpatched vulnerabilities, gaps through security misconfigurations, and finding and abusing valid credentials.

Internal pressure can come from insiders with potential bad intentions, but also, and especially in the multi-cloud, through the proliferation of shadow IT and shadow data across all these environments.

So why not solve all this with cloud-native tools?

It starts with a fundamental difference in philosophy, whereas Rubrik's enterprise data management is focused on a radically simple approach of providing abstractions for all workloads through an SLA policy-based system; a public cloud provider typically offers building blocks for you to create bespoke systems, not comprehensive enterprise solutions.

Not only does each public cloud have its backup tools, but even for a single cloud provider there are also differences between the type of workloads and how you set backup policies. Where backup data is stored is also different between workloads, which leads to challenges in making that data immutable to protect against ransomware etc.

If the tools are different, you'll also struggle to get a centralized view of your entire estate and understand the status and success of your data protection schemes, which leads to challenges around reporting, logging, and auditing.

A backup is only successful if we can use it to recover data. In the case of native tooling, this tends to be more basic, not very granular, and slow, with no insight into things like the potential reinfection of your environment after a malware infection.

Finally, a lack of enterprise-capable storage cost optimization, limited storage tiering, compression, and deduplication can lead to massive overspending. Multiply these challenges over multiple accounts, across multiple regions, across all hyperscalers, and include on-premises and SaaS workloads, and you end up with an unmanageable situation where probably no one can answer the question "Are we resilient?"

Building Cloud Cyber Resilience Without Compromising Your Cloud Bill

Building cloud cyber resilience without compromising your cloud bill requires strategic implementation of cloud backup solutions that balance both cost and security. By leveraging cloud-native APIs, serverless functions, and storage cost optimization strategies like tiering data to less expensive locations, and implementing data reduction techniques like incremental backups, compression, and deduplication, we can build out an architecture that delivers both resilience and keeps cost in check.

Cloud data still needs to be stored in an immutable manner to survive cyber threats and provide a quick and easy path for you to bounce back. By mandating immutability across your multi-cloud environment as standard you can avoid misconfigurations taking away your last chance of success in getting your organization operational again. Immutability can take on multiple forms, and we need to ensure you can survive both a data and account compromise.

By layering on cloud data security capabilities like anomaly detection and sensitive data discovery you can guarantee you are recovering a clean copy of your data and not risk reinfecting your cloud environment.

Operational savings can be achieved by abstracting away the underlying complexities of implementing individual cross-cloud workload properties but instead relying on centralized and unified SLA policies that easily translate your organization's operational requirements. A centralized solution will also provide the capability to report and monitor the success of the implemented policies without the need to explore individual areas across your cloud environments.

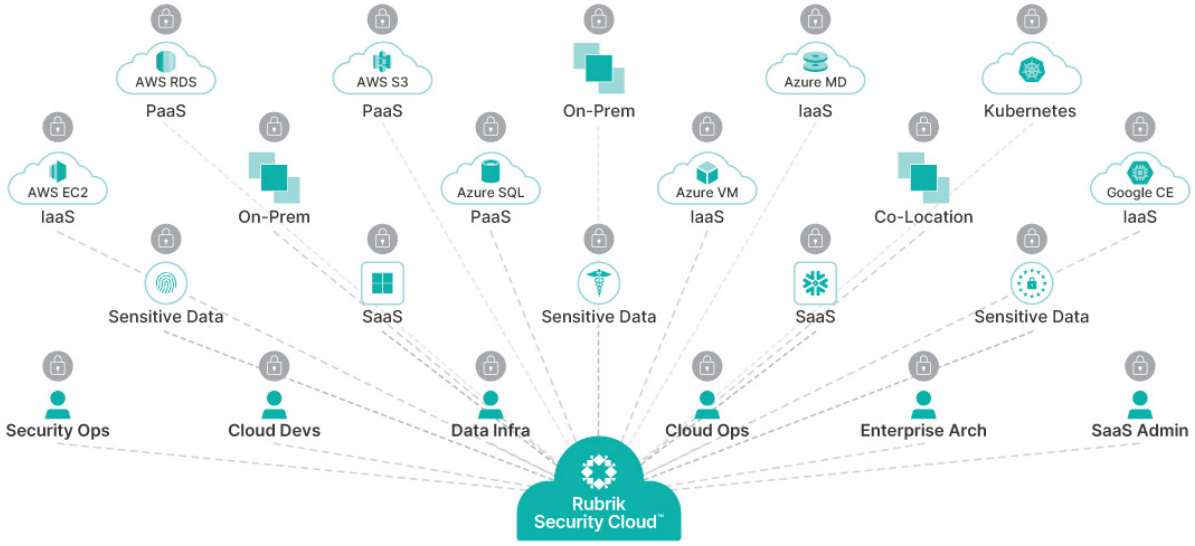

With Rubrik Security Cloud we help you first get visibility across all your data. What data do you have? Where does it live? Is it sensitive? Next, we provide the necessary context around your data. Who has access? Is it overexposed, redundant, or unprotected?

Once you understand your data, you can properly back it up. Which allows you to properly recover when needed, with cost-lowering benefits.

Written by:

Filip Verloy
Field CTO EMEA & APJ
Rubrik

Filip Verloy serves as Field CTO EMEA & APJ for Rubrik X. In that role, Filip engages and advises customers, partners, and the security industry at large, sharing his experience, insights, and strategies on data security. Before joining Rubrik, Verloy was the Global Field CTO at API security start-up Noname Security and has previously served at various IT vendors including Citrix, Dell, Riverbed, and VMware in roles ranging from Staff Architect to Solutions Executive supporting some of the largest and most complex customer environments. He has been in the IT industry for over 25 years, spanning the customer-, consulting-, and vendor side.


rubrik