KODIAK DATA

WHITE PAPER

Protecting Data in the Cloud: How Private Clouds Solve Security Issues

Introduction

Ninety-two percent[1] of companies are using public clouds in 2018, up from 89% in 2017. But while these companies are gaining all of the efficiencies and cost benefits of the cloud, organizations are concerned about maintaining the same level of security and control they have when running their applications on-premises.

This Kodiak Data white paper examines the security challenges of the public cloud, how security breaches are driving private cloud adoption, and how Kodiak Data can help your organization overcome these challenges with Kodiak Data's suite of private cloud solutions that provide you with the full benefits of running in the cloud along with higher performance and easier management.

Public Cloud Security Breaches

Data breaches are occuring at an alarming rate. While the media highlights a few of the more spectacular public breaches, multiple private companies[2],[3] have a far more complete list, and show that breaches are occurring several times per week.

One of the more interesting aspects of the breaches is the delay between when the breach occurred and when the affected users were notified of the breach. For example:

- The Yahoo breach[4] was first reported in December 2016, but had occurred in August 2013. Every single Yahoo account which was active in 2013, over 3 billion accounts, had been compromised for over three years before the affected people were notified.

- The Equifax breach[5] was first reported in September 2017 but had occurred in May 2017. 390 thousand people had Social Security numbers and credit card information exposed for over four months before they were notified.

- The Verizon and Orange S.A. breach[6] was first reported in June 2017 but had occurred starting in January 2017. Up to 14 million US customers, and millions more French customers, had their data exposed for six months before they were notified.

In all these cases, personally identifiable information was made public, and those affected didn't know about it for months or years.

There are other aspects about the Verizon and Orange S.A. breach that are even more disturbing. First is the geographic location of the data. In this breach, a third-party company was storing these records in Israel, completely outside of the physical and legal control of both telecommunications companies. Second, the third-party company was storing proprietary Verizon and Orange S.A. data on the same server without isolating and protecting this company confidential information. The result of this is that a single fault in the third-party company's security environment exposed multiple company's customer information at once.

On-premise environments, private clouds and public clouds can all have strong security, because they use the same security tools and methods. Even the US intelligence community has started using public clouds, but we should note that they required Amazon to use a separate facility and environment for their public cloud[7]. However, public clouds are shared workspaces managed by 'outsiders' and inherently lack the assurance of a dedicated security framework that can only come from the actual users themselves The inevitable requirement of public cloud environments is that they are outside of your control and accessible from anywhere.

Major Security Challenges Organizations Face Moving Data to a Public Cloud

As noted above, public cloud providers use the same tools as on-premise implementations to provide security for their environments, and do a very good job of it. Public cloud providers have the resources and the motivation to enable excellent security. However, their official policies state that security is largely the responsibility of the user of their service[8].

The entire business and technical model of public clouds requires full and complete remote access to the computing environment. In contrast to this, completely isolating production environments from public access in on-premise environments is not only easy, but has been the standard for decades. In many cases, this is done by placing them on networks which require physical access to the facility in order to use them, "air-gapping" them from outside contact.

This type of highly secure isolation is simply not possible in the public cloud which, by definition, requires public access, opening the door to assaults of every kind.

Outside of that level of physical security, there are additional security concerns around the public cloud:

Governmental regulatory concerns about the location of data[9]

The physical location of your data determines the legal jurisdiction of your data, specifically:

- Whose laws apply to your data? Is this settled, or will there be a legal argument about whose laws apply, or whether several jurisdiction's laws equally apply?

  Microsoft was sued by the US Department of Justice[10] concerning data whose indexes were stored in the United States, but whose contents were stored in Ireland. The question of whose laws apply, and whether a warrant from a US court applies to data stored in Ireland, is an open question[11].

- In each of those jurisdictions, what are the legal requirements for storage, retention, backup, protection, review and purging of personally identifiable information?

  If data is moved by the public cloud provider into a legal jurisdiction other than the one regulating your company, it is likely that the data will be subject to restrictions that you never considered. Does your data processing have a "follow the dawn" model? The City of Los Angeles had to put in a special clause with its contract[12] with Google to prevent just such movement. Does your contract with your cloud provider have such a clause? What are the extra costs of this?

Privacy concerns about storing company proprietary data off-site

This is different than the first issue, because it is a governance issue caused more by your senior management's concern around business risk and responsibility, than by governmental or other agency legal issues. Outsourcing business tasks to the public cloud means that the public cloud provider performs all of the day-to-day operations, and makes it more difficult for you to monitor and verify all aspects of business risk.

As shown by the Verizon and Orange S.A. breach, third-party companies may not take the full measures necessary to protect your data. They may put your data at risk without your knowledge, while you still retain responsibility for the consequences of the breach.

But the governmental and legal issues still apply. Not all governments are equally zealous about protecting the privacy of your data from the government itself. Your public cloud provider may be forced to reveal your data to governmental agencies anywhere in the world and then be ordered to not inform you of that release of data. For example, both Google and Microsoft have been ordered to release a customer's data including emails without informing the owner of those emails by the United States Stored Communications Act of 1986[10].

Providing high speed access to coordinated data to geographically diverse sites
Public cloud providers solve this problem with high cost data replication, including both the cost of the duplicate storage as well as ingress and egress charges. This can become extremely expensive.

How Private Clouds Solve These Security Issues

In contrast to public clouds, private clouds are entirely in your control. Whether they are on-premises or in a hosted facility, you have full governance on all aspects of the environment. Specifically:

Data access: With private clouds you decide whether you require an air-gap to protect your data, and what level of controls you place on connection to your entire environment. There is no requirement to allow external access of any kind, up to and including Tempest certification.

Data location: Private clouds allows you to can evaluate the legal jurisdiction and implications of the data location as simply part of the decision process of implementing your business strategy. And the data will not be moved to another location without your full knowledge and prior approval.

Data monitoring: Any data breach, including those by governmental officials using their investigative authority, will be known to you immediately.

Cost of replication: In a private cloud environment, you control the amount of duplicate storage, with no ingress or egress charges.

Bottom line, private clouds solve the security issues inherent in public clouds.

How Kodiak Data Provides the Best Solution to These Security Issues

Using industry standard services and community compliant environments, Kodiak Data provides private cloud infrastructure that can be accessed as a hosted service or deployed on site or in co-located data centers. You can securely run all of your cloud-based applications, such as ClickHouse, CloudEra, ELK, Hadoop, HortonWorks, Kafka, Kubernetes, etc., on the Kodiak Data private cloud, completely avoiding all of the issues discussed above.

In addition, very large instances of Oracle and Microsoft SQL Server can be run in this environment to provide cloud efficiencies for your traditional databases. Designed for big data applications, our private cloud infrastructure can grow dynamically and non-disruptively to any level needed. By reducing the number of objects to manage, you reduce the number of entry points for security breaches. The Kodiak Data private cloud also does not provide any access to the applications or data and, subsequently, does not present an entry point for breaches. Internal components such as memcached are not exposed; therefore, they are not vulnerable to DDoS attacks. Your existing protection schemes will work without change when running in our environment.

Kodiak Data can also lower your storage costs with another software feature, EdgeCache, which makes data available via caching, and eliminates the need to replicate data while still maintaining local speed bandwidth and latency to that data at remote sites. Overall, with our solution, you reduce your storage capacity, your data stays safely in the single data center and legal jurisdiction that you specify, and your remote sites have effectively local access to that data.

Conclusion

Public clouds provide many benefits to organizations, but they also present security challenges. The use of private clouds, whether on-premises or in a hosted facility, can deliver the same benefits without the security challenges.

The Kodiak Data private cloud speeds up data- and analytics-heavy workloads by orders of magnitude, reducing the cost to both build and operate to a fraction of the cost of public clouds with comparable requirements, and is not subject to the security issues of a public cloud. We provide the fastest, most cost effective, and easiest-to-manage private cloud environment available to enable you to achieve your business goals more quickly at reduced cost.

For more information,
please visit https://www.kodiakdata.com.

[1] https://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2018-state-cloud-survey
[2] https://www.privacyrights.org/data-breaches
[3] https://www.idtheftcenter.org/2017-data-breaches
[4] https://www.oath.com/press/yahoo-provides-notice-to-additional-users-affected-by-previously/
[5] https://www.equifaxsecurity2017.com/consumer-notice/
[6] https://www.upguard.com/breaches/verizon-cloud-leak
[7] http://www.govexec.com/magazine/features/2014/07/daring-deal/88207/
[8] https://aws.amazon.com/compliance/shared-responsibility-model/
[9] https://searchcloudsecurity.techtarget.com/tip/Cloud-computing-legal-issues-data-location
[10] http://www.scotusblog.com/case-files/cases/united-states-v-microsoft-corp/
https://www.lawfareblog.com/primer-microsoft-ireland-supreme-courts-extraterritorial-warrant-case
[11] http://www.bbc.com/news/technology-27191500, comments by Mina Andreeva, European Commission
[12] http://clkrep.lacity.org/onlinedocs/2009/09-1714_rpt_cao_10-7-09.pdf, Appendix J.1 Section 1.7