parametrix
insurance

Managing Cloud Outage Risk

Cloud Downtime in 2022

Contents

A Note from Jonathan Hatzor

Downtime is an inevitable reality for almost every business. As companies increasingly rely on cloud providers for day-to-day operations, they are exposed to the real risk of downtime. It can occur any day, is entirely beyond their control, and can cause significant disruption.

Dependence on IT infrastructure and the digital value chain was illustrated in 2022 through events such as the Southwest Airlines crew scheduling system breakdown. It resulted in an extraordinary loss of up to $1.2 billion in December. The event shows how every company with a computer file stored off site needs to pay attention to the risk of downtime.

Fortunately, risk managers now have access to flexible and responsive insurance products that can effectively transfer cloud downtime exposures, including non-physical-damage interruption risk. At Parametrix, we have developed coverage that protects companies against all downtime-related losses, including lost revenues, customer compensation, SLA liabilities, reduced productivity, failed orders, and reputational harm.

Our powerful combination of technology and skills allows us to deliver solutions to this enormous, growing, and already leading risk of the new digital age, one which conventional cyber insurance policies do not cover effectively. In the following pages we reveal the downtime risk experienced by businesses in 2022, based on the findings of our comprehensive cloud monitoring technology.

Jonathan Hatzor,
Co-founder and Chief Executive Officer, Parametrix Insurance

Cloud Outage Risk

Businesses world-wide are increasingly reliant on the cloud — comprising computers and software services owned by others — for their day to day functionality. In a recent survey by Parametrix, **95% of corporate decision-makers said their business is dependent on the cloud. Nearly half said the cloud is a mission-critical service. Because of the cloud's enormous new importance in the digital supply chain, the cloud economy is growing at an unprecedented rate**. Worldwide end-user spending on public cloud services is expected to reach $591.8 billion in 2023, up two points from 18.8% growth in 2022, according to the latest forecast from Gartner. This breakneck growth is driven by the many benefits that the cloud offers, including flexibility, accessibility, and scale.

The cloud is reliable, until it goes down. When it does so, some businesses grind to a halt, often for many hours. Various problems can lead to a range of disruptions which cause impacts that range from minimal to critical. The more severe and prolonged the downtime event, the greater the cost impact. At the extreme end, cloud downtime can be very costly indeed. In the Parametrix survey of corporate decision-makers, more than half said cloud downtime would cost them $100,000 per hour or more. More than 5% who said it would cost their employers in excess of a million dollars per hour.

As exposure to this developing risk increases, companies more often identify cloud outages as a key operational risk in financial reports and business continuity plans. They have increased expenditure on redundancy and disaster recovery technologies and processes intended to minimize their Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO). They increasingly recognize cloud downtime as a risk that must be managed.

Concentration

The supply of cloud services is highly concentrated, with about two-thirds of global supply provided by three main companies: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). Naturally these companies typically report only major disruptions to their services, yet still they reported hundreds of performance interruptions per year in 2021 and 2022 — an average of 25 per month. In other words, even by their own account, the cloud goes down almost every day.

When, Not If

Reliance on the cloud is set only to increase in importance across the economy. Providers will continue to improve reliability, but mistakes and shortfalls will still lead to downtime events, and therefore to business interruptions and very real financial losses (if not to the physical damage covered by most insurance policies). Companies must recognize, mitigate, and transfer this risk, which we enumerate more clearly in the balance of this report.

The cloud is reliable, but it isn't perfect. Its imperfection is driven in part by a tradeoff between innovation and mistakes that result in errors. This tradeoff means that **100% uptime cannot be achieved when technology continues to improve, as it does almost constantly**.

The Register

Day 7 of the great Atlassian outage: IT giant still struggling to restore access

Majority of affected users still wondering where their data went. icon Richard Speed. Mon 11 Apr 2022 // 12:30 UTC.

11 Apr 2022

Mashable

Half the internet died while you were sleeping. Here's what ...

Cloudflare (the network, proxy & security provider for all these companies) is facing a global outage. 12:11 AM · Jun 21, 2022.

21 Jun 2022

MyBroadband

Major AWS outage in South Africa

Amazon Web Services (AWS) appears to be having an outage causing headaches for Elastic Compute Cloud (EC2) customers in its South African...

14 Sept 2022

DatacenterDynamics

AWS US-East-2 experiences issues "impacting Internet ...

AWS US-East-2 experiences issues "impacting Internet connectivity". But far from a full scale outage. December 06, 2022. By: Sebastian Moss. Have your say.

6 Dec 2022

CNN

Massive Southwest Airlines disruption leaves customers stranded and call centers swamped

Long lines at Southwest counters Monday, Dec. 26, 2022, at Raleigh- · Southwest canceled about two-thirds of its flights. See how travelers are...

26 Dec 2022

BankInfoSecurity

Microsoft Experiences Second Major Cloud Outage in 2 Weeks

7, 2023 15:53 UTC: This story has been updated to include Microsoft's assessment that the incident has been almost fully remediated. Business...

1 month ago

Understanding Cloud Risk

Leading companies and their risk managers have begun to map and understand the potential impact of cloud outages on their business. Risk managers and other business leaders are rightfully concerned about the potential costly impact of a cloud outage on their business. In a survey of 324 US-based corporate decision-makers conducted for Parametrix:

- 95% said their company is **reliant** on one or more third party cloud-based services
- 60% said they are **very concerned** about cloud downtime
- 82% said their organization and/or product is **dependent** on their availability of the cloud
- 41% indicated they are **more concerned** about downtime this year than last year
- 31% said eight hours of cloud downtime during business hours would be **catastrophic**
- 36% said eight hours of cloud downtime would be a **non-recoverable**, **uninsured** loss
- 52% said a downtime event **upsets customers** and **increases churn**
- 50% said it causes **lost revenue** and sales
- 52% said **manufacturing or office work halts** during cloud downtime
- 60% of decision-makers with revenues of $500M and above said downtime events led to **direct revenue loss**
- 23% estimated downtime costs their organization **$500,000 or more per hour**

An increasing number of companies now reports their concerns over cloud downtime risk in formal stock exchange documents such as annual 10K filings. Examples include:

salesforce

> ***Any interruptions or delays in services from third parties***, *including data center hosting facilities, cloud computing platform providers and other hardware and software vendors, or from our inability to adequately plan for and manage service interruptions or infrastructure capacity requirements,* ***could impair the delivery of our services and harm our business***.

DOORDASH

> *We currently host our platform and support our operations on a* ***single data center*** *... We have experienced, and expect that in the future we will experience, interruptions, delays, and outages in service and availability from time to time ... Any of the above circumstances or events may harm our reputation and brand, reduce the availability or usage of our platform,* ***lead to a significant short-term loss of revenue***, *increase our costs, and impair our ability to attract new users, any of which could adversely affect our business, financial condition, and results of operations.*

NETFLIX

We rely on Amazon Web Services ... and any ***disruption of or interference with our use of the Amazon Web Services operation would impact our operations and our business would be adversely impacted*** *... Currently, we run the vast majority of our computing on AWS.*

We primarily rely on Amazon Web Services *to deliver our offerings to users on our platform ... AWS facilities are vulnerable to damage or interruption from natural disasters, cybersecurity attacks, terrorist attacks, power outages and similar events or acts of misconduct .... Any* ***negative publicity arising from these disruptions could harm our reputation and brand*** *and may adversely affect the usage of our offerings.*

Cyber attacks or data breaches on our customers' or third-party providers' networks, or in cloud-based services provided to, by, or enabled by us, ***could result in claims of liability against us, give rise to legal and/or regulatory action****, damage our reputation or otherwise materially harm our business ... [and] could, in each case, result in claims of liability against us, damage our reputation or otherwise materially harm our business.*

Financial Impacts

Cloud outages are inevitable and potentially very expensive. No matter how much redundancy engineers build into cloud-based systems, they will continue to occur for reasons including connectivity errors, equipment failure, cyberattacks, and — most commonly — human error. Outages that last only minutes cause inconvenience, but those that last for hours can have disastrous consequences for businesses.

According to Information Technology Intelligence Consulting (ITIC), **40% of enterprises said a single hour of downtime could cost between $1 million and $5 million, exclusive of any legal fees, fines, or penalties**. In a Gartner survey, 98% of companies stated the cost of IT downtime ranged from $100,000 to $540,000 per hour.

Outages present a major threat to operational resilience. They affect companies directly, but also their customers and their customers' customers. Five primary risk areas may be impacted severely by outages: **financial, reputational, legal, operational, and fulfillment**.

Financial risks include lost revenues and workaround expenses. A cloud outage can shut down critical sales channels at any time, preventing customers from initiating a purchase. Some may come back another time, but others will seek an alternative and never try again. Boston-based e-retailer Wayfair achieved $14.1 billion in online sales in 2020, roughly $1.61 million every hour. Four hours of downtime could cost the firm $6.44 million (less if it happens in the middle of the night, and a lot more if it's on Cyber Monday).

Reputational risks include brand tarnishing due to non-performance arising from cloud downtime which created disappointed customers. Social media is fertile ground for sharing negative sentiment, and drives customers to look elsewhere for available services. Nearly 37% of

small businesses have reported that they lost customers due to downtime.

Legal risks can arise from downtime when contractual obligations are missed, when shareholders note underperformance due to losses arising from outage, when regulators assess customer service as falling short of expectations due to a failure to address downtime, and in a host of other areas.

Operational risks include lost productivity. Since many companies rely on the public cloud for almost all file use and management, internal and external communications, development, and other key operational functions, lost productivity can be a major cloud-downtime impact. According to the Society for Human Resource Management, salaries account for 18% to 52% of any business' operating budget. For a company with a $500 million operational budget and a 30% salary spend, a workday lost to downtime costs $575,000 in salaries paid to employees who can't work.

Fulfillment risks could include missed Service Level Agreement thresholds. Many businesses are contractually obliged under such deals to provide specified services within a specific period, but may find fulfillment impossible due to cloud outage. Companies in the Software as a Service and FinTech sectors are particularly at risk, since every second of downtime may have enormous fourth-party customer consequences.

A single downtime event may impact many companies simultaneously. In January 2023, MS Azure experienced an outage event which affected Azure resources in multiple regions. The incident lasted for two hours and 24 minutes, according to the proprietary Parametrix Cloud Monitoring System, which identified the event immediately. Microsoft Azure is the world's second-ranked cloud service provider. The downtime affected thousands of businesses and tens of thousands of users. Most companies impacted were entirely uninsured.

Figure 1: MS Azure Outage, 25 January 2023, as recorded by Parametrix Cloud Monitoring System

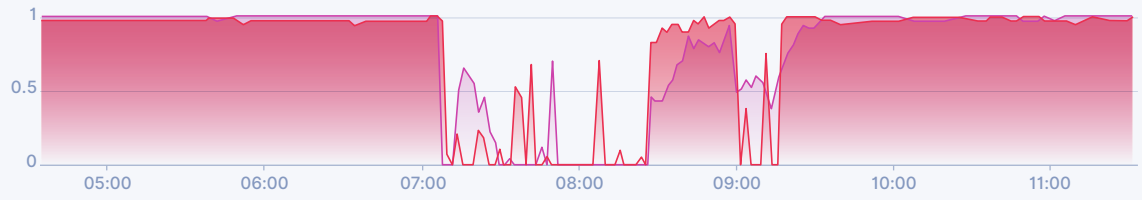

Cloud Outage in 2022: the Parametrix Findings

Superior Monitoring by Parametrix

Parametrix uses proprietary technology to monitor the availability and performance of these top cloud service providers worldwide. **We conducted 36.4 billion availability and functionality tests on more than 250 cloud data centers in 2022.** The tests allow us to assess the potential business interruption and associated risks related to each provider, supporting better comparison, indexing, modeling, and pricing of downtime events.

We measure the full outage periods that the providers report, but also impaired services and other disruptions which are not well reported by cloud service providers. We do so in great detail. We then assess the impact of each interruption or degradation from a cloud user's perspective. That allows us to structure risk transfer products which react to the genuine disruptions which lead to specific businesses' actual dependencies. **We monitored nearly 1,200 performance interruptions and disruptions in 2022, four times the number reported by the providers themselves.**

The analysis of the following cloud outage events is focussed on the three main cloud providers. Together these providers provide nearly two thirds of the world's cloud services, according to DgtInfra.

Between January 1st 2022 and December 31st 2022 Parametrix Cloud Monitoring System (PCMS) conducted **36.4 billion** availability and functionality tests on more than **250 cloud data centers** (known as Availability Zones, or AZs) around the world[1], encompassing all of the **top three cloud service providers**. Compiling the data collected throughout 2022 enabled us to deliver a detailed picture of the cloud performance during the year.

The Parametrix Cloud Monitoring System (PCMS) monitors each cloud provider at the levels of service, region, and AZ, as is relevant to each specific provider. This report measures and summarizes the number of individual impacted service units (or impacted units) for the monitored events. Each of these impacted units reflects an individual instance of impact to a specific cloud service in an individual region. For example, an event affecting both the Database and Compute services in the 'us-east-2' region would be considered two separate impacted units of one event. PCMS also considers dependencies between services.

Cloud users rely on the cloud for various services supplied from different regions, depending on their needs. For example, one user might use Storage in the us-central1 region, while another may use the Compute and Functions services in the same region. In this case, an outage event impacting Storage alone would not impact the second user. That makes the granularity of monitoring and measurement methodology critical to the assessment of event impacts.

In 2022, Parametrix identified a total of 1190 performance disruptions across this cloud landscape. Some 41.4% of these events, 492 in total, were classified by Parametrix as critical. However, this simple count of events does not indicate the level of performance impact, nor

[1] *Excluding China and government data centers*

the availability of services of each cloud provider during the events monitored. From a risk perspective, further factors must be considered when evaluating provider-level performance. These may include:

- total duration
- the duration of each impacted unit within the event
- the mission-critical nature of the service to its users
- the number of Availability Zones affected
- and multiple additional variables.

The fourth quarter of the year was the most stable quarter. In terms of geographical impact and timing, about a third of events impacted US regions. The balance of the critical events measured were split roughly evenly between Europe, Asia, and the rest of the world.

Figure 2: Geographical split of critical outage events as detected by PCMS

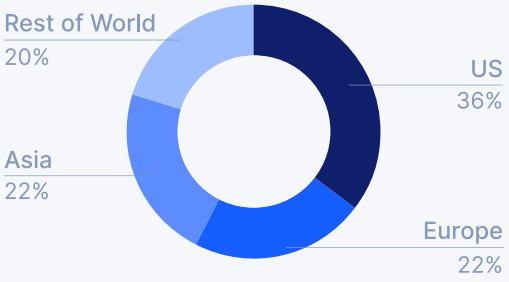

Causes of Disruptions

Various root causes lead to disruptions and critical events. Downtime occurs due to a mix of software, hardware, and infrastructure issues, despite cloud providers' heavy investment in data-center resilience, which is a top priority for them.

Reported root causes for events suffered by AWS include bugs triggered by increased traffic and power outages. Azure and GCP reported causes including issues with physical infrastructure, insufficient resources due to increased traffic or recent rollout, and misconfiguration during maintenance. The most common reported root cause of outage events in 2022 was human error, which includes misconfiguration and faulty maintenance activity.

Figure 3: Root causes and average durations of critical outage events in 2022

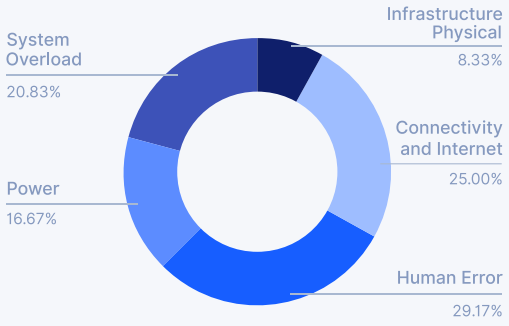

Parametrix Stability Index

Each year Parametrix releases the Parametrix Cloud Stability Index (PCSI), which is based on data from its proprietary PCMS. This unique index is based on the overall performance of cloud providers' regions during the year. It considers, among other factors, the number of events per region and service, their relative severity, the duration, and the impact of each event. The index measures relative performance during the relevant period only, and does not reflect prior experience and is not a projection of performance and risk. Scores are calculated annually based on historical data for internal modeling and product development purposes.

The Index values are classified according to five levels of stability, with those rated "A" designating the best performing regions (with a score of between 0 and 20), and those rated "E" the worst performing (scoring 80 to 100). Index results inform a long-term predictive model of future downtime risk by region. PCSI supports the risk assessment and consulting services Parametrix supplies to clients and partners about the resilience of their cloud architecture. The bubble graph in Figure 4 illustrates all providers' US regions. The size of each bubble indicates usage, and the color its score, from A (lightest) to E (darkest).

Figure 4: US regions usage level and Parametrix Stability Index in 2022

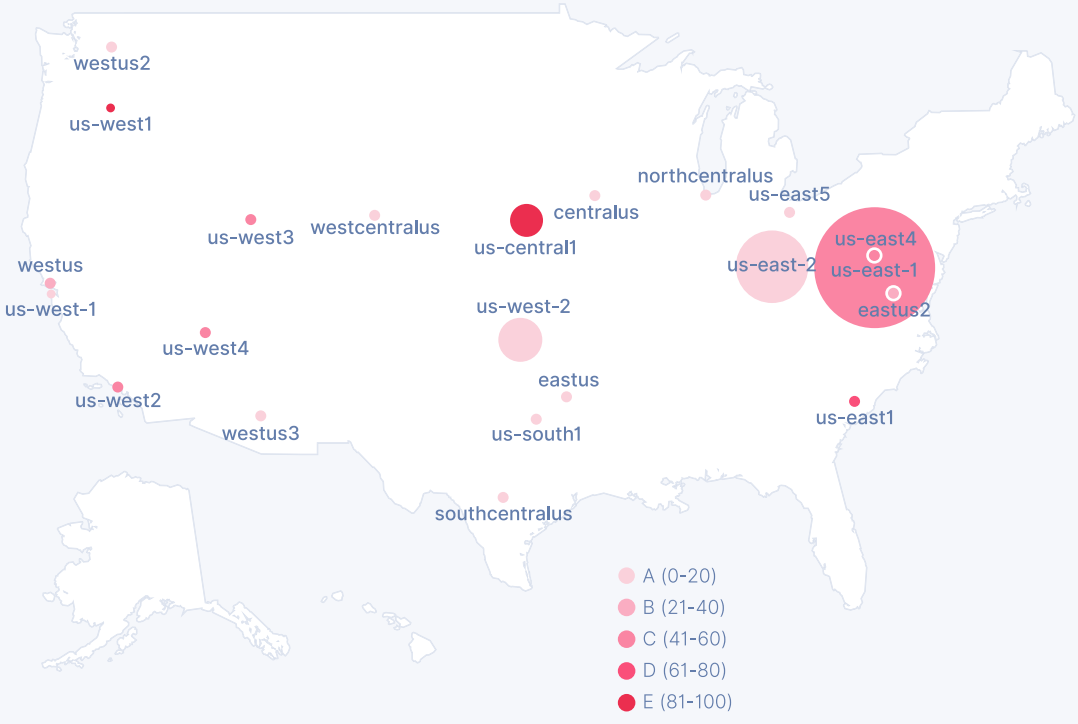

Figure 5 illustrates the percentage of each provider's regions which fall into each stability index category.

Figure 5: Market leaders' percentage of regions by Parametrix Stability Index in 2022

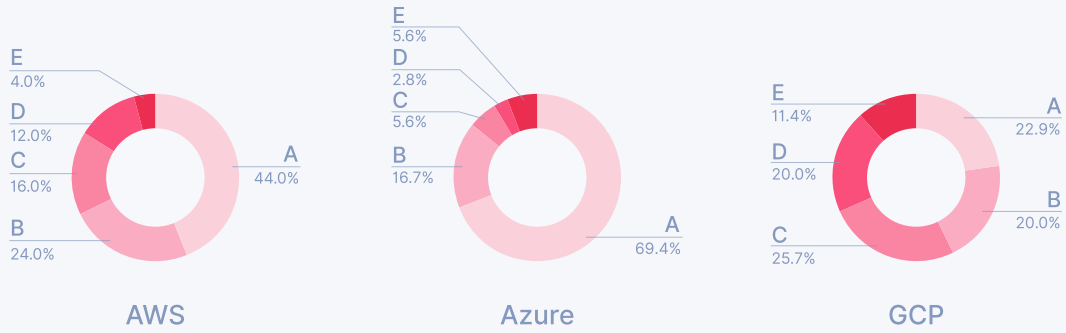

Figure 6 illustrates the number of each cloud provider's services affected during 2022, both at the critical and warning levels.

Figure 6: Number of services affected in critical and warning severity levels for each cloud service provider

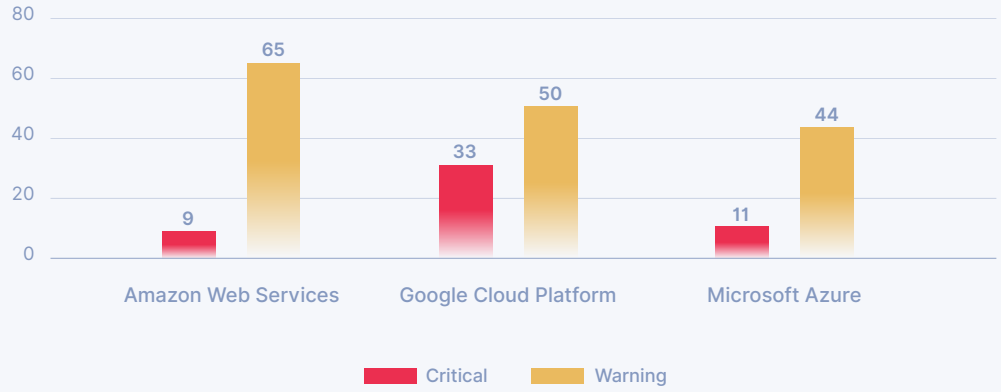

PCMS identified 240 services which suffered a performance interruption in 2022. Of those 53 were critically impacted across the major providers. Some events affected multiple services. For example, AWS Autoscaling Service, which enables users to automatically increase or decrease their compute power to meet business needs, was affected during five different critical events.

Parametrix Event Analysis

The following event analysis describes one of the many Critical Outage events that occurred during 2022 and was detected and monitored by PCMS. It reveals the resolution of the information collected during our monitoring of one of the more significant events of the year.

Event Analysis

Service provider: Azure

Region: East US 2

Start date and time: April 8 2022, 12:17 PM UTC

End date and time: April 9 2022, 02:40 AM UTC

Total event duration: 14:23 (Some services were not affected for the entire period and their impacted units were of shorter duration)

Main affected services: Azure Virtual Machines (VMs), Azure Autoscaling

Root cause: Insufficient resources to support increased traffic

Description and impact: PCMS detected an increase in errors which prevented users of Azure VMs in EAST US 2 to create, delete, or update VMs in all availability zones. The Autoscaling service was also affected severely in two of three AZs, affecting users who tried to adjust or scale their existing computing power to their business needs.

Figures 7, 8, and 9 illusrate the success rate of impacted services, the Y axis represents the success rate, as detected by PCMS.

Figure 7: Success rate of Autoscaling service in Azure East US 2 region, Zone 3

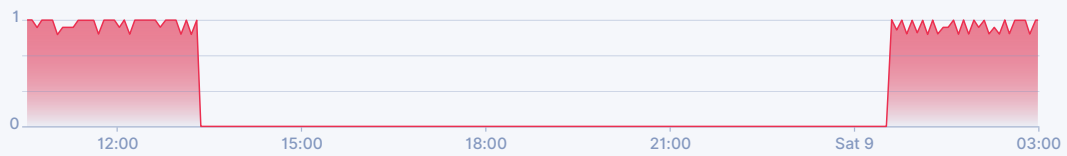

Figure 8: Success rate of Autoscaling service in Azure East US 2 region, Zone 2

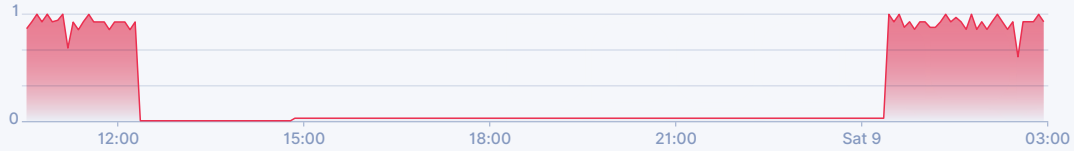

Figure 9: Success rate of Virtual Machines service in the Azure East US 2 region

parametrix insurance

Parametrix is a pioneer in the monitoring and modeling of cloud providers and cloud-based services. Our platform automates risk assessment, pricing, and modeling of these sophisticated platforms. Parametrix continuously monitors a variety of third-party IT services across the globe to collect granular data on service interruptions. We generate real-time alerts on outage events, and provide access to valuable data sets. As a Managing General Agent and Lloyd's Coverholder, Parametrix underwrites parametric insurance against outages of the cloud, CDNs, or e-commerce platforms, backed by major A-rated global insurers. Parametrix operates from offices in New York, Tel Aviv, London, and Munich.

www.parametrixinsurance.com

Parametrix Insurance

info@parametrixinsurance.com