



PROACTIVE OBSERVABILITY

Protecting Revenue Through SLA Monitoring

Learn To Manage Service Level
Agreements With Data You Can Trust



“You can’t enforce agreements without monitoring SaaS applications, and with Catchpoint we’ve been able to be refunded millions.

Maira Zarate, Application Monitoring Engineer, Autodesk

Table of Contents

- 3** Introduction
- 4** The changing nature of IT
- 4** The importance of service level agreements
- 5** SLA observability
- 6** SLA metrics
- 7** Example SLA statements
- 7** External and internal providers
- 8** Improve accountability with the right data
- 8** How to protect revenue streams with service level management
- 9** Ensure your SLAs are broad and robust
- 10** Observe your data in the right way
- 11** Where you observe from matters
- 12** The biggest SLA trap
- 13** The need for a reliability strategy
- 13** How to implement a multi-vendor reliability strategy
- 14** Why to consider a backup reliability strategy
- 14** What is a rock-solid observability strategy
- 15** Conclusion

Introduction

Virtually every enterprise today is adopting new digital transformation strategies to deliver value to its customers. The enterprise increasingly relies on digital to interact, communicate, sell to, and service its consumers, and increase the efficiency of business operations.

Meanwhile, IT is adopting agile development processes to build and deliver digital services. This means businesses are increasingly dependent on cloud providers and other third-party vendors for key components that are critical to the functioning and delivery of their services.

The concept of digital transformation has become a catch-all term and can mean different things to different companies. [CIO recently interviewed](#) a number of businesses about their digital vision. The CIOs defined digital in a variety of ways, including:

“**Digital is all about how we reach the customer. Traditionally, Western Union has been known as a cash business; our goal now is to digitize the customer experience wherever we can. We measure the success of our digital transformation by tracking the channels our customers choose to use: web, phone, or an agent location. Our transition to digital is all about giving the customer convenience, simplicity, and options.**

Sheri Rhodes, CTO, Western Union formerly;
now CIO at Workday

“**For Lenovo, digital ranges from basic process optimization, to using technology to unlock new business models, to creating new products, and delivering more empowering customer experiences. We believe in the increasing importance of artificial intelligence for businesses going forward and have been using the term ‘Intelligent Transformation’ — applying technology, especially AI, in all of the areas of the business to tap into the exploding amounts of data that are becoming more available throughout the enterprise and ecosystem.**

Arthur Hu, CIO, Lenovo

The changing nature of IT

The number of companies providing a digital service to enterprises has skyrocketed in the last decade, with the widespread shift to cloud services. An ever-increasing reliance on cloud providers to deliver digital services is dramatically changing the role of enterprise IT beyond architecting, developing, and monitoring performance, to adopting and governing the cloud.

Enterprises have to figure out which infrastructure types best suit their individual workloads and what the overall profile of infrastructure types should look like. The good news is that the wide range of cloud providers and third-party services available enables businesses to determine the best infrastructure for their specific use cases.

Private cloud solutions that offer dedicated resources are **growing the fastest** in share of compute resources, suggesting that the adoption of Infrastructure as a Service (IaaS) and Software as a Service (SaaS) resources will continue to grow in the future.

The importance of service level agreements

For years, procurement and legal departments have ensured that contracts with service providers contain strict clauses on Service Level Agreements (SLA), including any penalties that will incur if SLAs are broken. As digital becomes an increasing component of everyday business, the importance of SLAs in maintaining good service and governing the cloud is becoming ever more important.

Properly set and enforced SLAs provide the company consuming the service (the customer) with objective grading criteria and protection from the business impact of poor service. The service provider, meanwhile, gains from the opportunity to set appropriate expectations for how its service will be judged and, because it is being held accountable, is incentivized to improve quality of service.

Major service outages

Numerous major service providers have experienced significant outages, demonstrating the importance of having robust SLAs in place. These include:

Slack – On 22nd February, 2022, **Slack users were unable to access** the communication tool for a significant portion of the morning. It was Slack's **second major outage** within a six month period.

Amazon Web Services – December 2021 saw **a trifecta of AWS outages**, which not only took down Amazon, but a slew of other dependent services such as Kindle, Netflix, and Zoom.

Google Cloud – A similar downstream impact was felt on November 16, 2021, when **Google Cloud suffered a serious outage**, impacting many companies which rely on GCP for hosting, from CNET to Spotify.

Telia – Meanwhile, Telia, a major backbone carrier in Europe, **suffered from a networking routing issue** on October 7, 2021, causing yet another ripple effect, impacting Cloudflare, Equinix Metal, NS1, and others.

Facebook – October saw perhaps the biggest outage of last year when Facebook and its associated apps (from Instagram to WhatsApp) **went down simultaneously around the world** for an extended period.

Akamai – The summer of 2021 saw a slew of outages and degradations for some of the world's most widely used CDNs, with **Akamai experiencing issues** on July 22, 2021 and August 31, 2021.



Preventing SLA breaches

Having an SLA signed and on file alone is not sufficient protection. The customer must ensure the provider is meeting its SLA, and the provider must ensure it does not breach the SLA.

Additionally, almost all SLAs require the customer of the service to file a breach request to trigger any penalties defined therein. In other words, a service provider will not provide you the credits that the SLA states you are owed unless you specifically request and provide proof of a problem.

When Microsoft 365 had a two-hour outage in April 2019, they breached their SLA with Catchpoint. Even though the outage was publicly known and verified, Catchpoint still had to file an SLA request, which a special Microsoft team then validated and verified, ultimately agreeing to issue a credit to our account.

When we speak to companies before they deploy our observability solution, we have frequently discovered that many of them have failed to gather any penalties from their vendors in the event of a breach. This leads to finger pointing and a deterioration of the relationship.

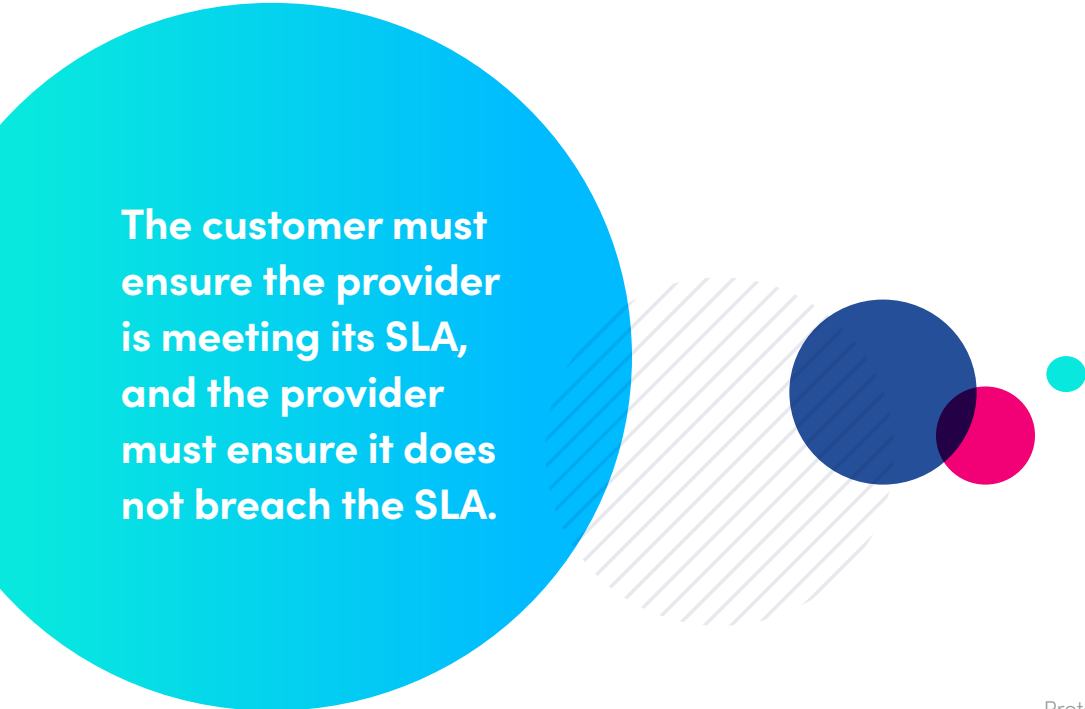
SLA observability

SLA observability is an often-overlooked aspect of properly managing SLAs. Typically, a Service Level Agreement is based on monthly data. Therefore, monitoring the agreement on a monthly basis is essential for proactive detection of breaches.

Observing your service providers and using the right enforcement strategy can:

- Ensure the provider focuses on delivering you the level of service agreed on.
- Improve accountability.
- Help recoup any costs of failure.
- Get you out of a bad relationship, if and when necessary.

Let's see what customers of cloud services can do to measure the performance and manage the SLAs of their various cloud providers, and what traps and pitfalls to avoid along the way.



The customer must ensure the provider is meeting its SLA, and the provider must ensure it does not breach the SLA.

SLA metrics

The term SLA is widely used and has thereby become an umbrella term. A company cannot actually observe an SLA, which is really just a document that outlines the terms of service. A business can only enforce an SLA by observing the metrics referenced within it.

A metric referenced by the SLA, which is a quantitative measure of the level of service, is called a Service Level Indicator (SLI). The SLI is measured in relation to a service level objective (SLO), a goal that cannot be breached. The SLO provides a value or range of values considered acceptable for the SLI. There is generally an upper or lower-bound limit. For example, DNS resolution time may be reported as not taking over 100 ms.

Almost all SLAs for digital services have an availability metric as an SLI. Availability is a measure of time within a given time-period, typically the calendar month, for

which the service(s) was reachable and functioning as expected. Therefore, the availability metric measures the reachability of the service from outside the provider's own infrastructure, and checks that the code behind the service is performing the function as expected.

If you are using Twilio as an SMS delivery service, for instance, the availability metric would measure the fact that Twilio's API is reachable, and that it properly responded to the API request to send an SMS, in addition to whether the SMS was delivered or not.

Any SLA will outline how the SLI will be measured, the length of time or number of measurements that must be outside the range, and if there are any consequences if the agreement is breached. If there are no consequences, there is no SLA.

Service Level Indicator (SLI)	Metric That Is Measured	Availability
Service Level Objective (SLO)	Acceptable range of values for the SLI	<ul style="list-style-type: none">• 99.99% – 100%• 99.5% – 99.99%• < 99.5%
Service Level Agreement (SLA)	Legal document or contract with end users defining consequences if SLO is not met	<ul style="list-style-type: none">• 20% credit for breaching 99.99%• 50% credit for breaching 99.5%

Availability	Downtime Per Month
99.0%	7.3 hours
99.5%	3.65 hours
99.9%	43.83 minutes
99.95%	21.91 minutes
99.99%	4.38 minutes
99.999%	26.3 seconds



Example SLA statements

“The service provider will use commercially reasonable efforts to maintain a Network Connection rate of **99.9% per month** (that is, the aggregate monthly network failure does not exceed 44 minutes).”

“The service provider will make its dashboard available with a monthly uptime percentage of **at least 99.5% during any monthly billing cycle** (the “service commitment”). In the event the vendor dashboard does not meet the service commitment, you will be eligible to receive a service credit as described below. Downtime is defined as any period of time during which our service is not available for customer use due to a network, hardware, or software failure within the service provider’s data center, hardware, or vendor itself. Downtime does not include unavailability of the vendor’s service due to customer-specific issues such as network, hardware, or software problems at the customer site. Downtime will always be rounded up to the nearest minute, and expressed in terms of minutes of downtime.”

“‘System availability’ means **the percentage of total time** during which the SaaS service is available to customer less the scheduled maintenance downtime, which should be less than 8 hours per month and be performed during the hours of (i) midnight Friday to 7:00 a.m. Saturday US eastern time or (ii) midnight Saturday to 7:00 a.m. Sunday US eastern time.”

“Supplier will provide **99% system availability over one-month periods**, excluding any system maintenance or force majeure events (as defined below) that result in the system not being available to any customer user, as measured and monitored from supplier’s facilities.”

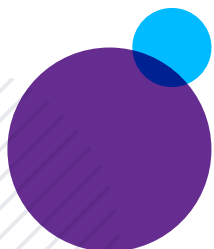
External and internal providers

As mentioned earlier, an SLA exists between the customer and the provider of services. The provider can be an external or internal provider, for example, another group or division within the company. External providers include DNS services, content delivery networks (CDNs), managed service providers, cloud compute, hosting providers, validating services, translation services, API platforms, DDoS protection, Bot Protection, Identity and Access Management (IAM), productivity software, communication and collaboration suites, and more.

Internal SLAs are no less important than external ones. In large enterprises, different groups have different budgets and priorities. However, there are often dependencies on the services, applications, or infrastructure that can result in one division impacting another, and causing a domino effect of outages and a negative impact to business. Internal SLAs help the various enterprise groups hold one another accountable, and ensure that the quality of service delivery takes priority without finger pointing and political infights.

It doesn’t matter whether the provider is internal or external; if the application or service being provided is business critical, an SLO and SLA should be set.

However, setting an SLO and publishing an agreement is not enough. You also need to accurately measure the service to ensure that the SLOs are being met.



Improve accountability with the right data

The overall goal of the **SLA is to improve the accountability** of the provider and ensure that the customer paying for the service can mitigate any risks involved by shifting service to a different provider when necessary.

To improve accountability, it is critical that the right data on the service is collected. It must accurately represent the expectations of the SLA and what the service provider has within their control. Without accurate data, the relationship can quickly turn into a finger pointing exercise in which each party feels that they are right. “Yes, the performance wasn’t what was expected, the customer is right.” “No, the SLA wasn’t breached, the provider is right.” Looking at things from a neutral standpoint, however, often reveals that there is no one right answer.

Secondly, the provider can exercise control over the quality of what is delivered only up to a certain point on the Internet. **A provider of a digital service cannot be responsible for what is happening outside the realm of its control.**

For example, if there is a major fiber cut in Virginia, that impacts the Internet for most of the Northeast; a content delivery network (CDN) provider will be impacted, but the provider will be unable to do much about the cut and the outage is not their responsibility. Therefore, the first step in bringing accountability and governance to SLAs is to collect objective and accurate measurements.

Some service providers might issue reports on how they are performing in relation to SLAs, but verifying this with your own observability efforts will increase confidence in the level of performance being received. It ensures trust on both sides, and means that you won’t feel cheated if and when your organization experiences business pain while your provider’s data appears to show that everything was normal. Where the metrics are observed from matters.

How to protect revenue streams with service level management

Enterprises started to monitor SLA metrics in the late 1990s. At the time, businesses focused on monitoring websites, ad-serving systems, or HTTP URLs for their primary providers for hosting companies, ISPs and CDNs.

By contrast, as we have discussed, the typical enterprise today relies much more heavily on external providers of services, which deliver web pages and static image files, as well as DNS, APIs and other important client server protocols.

Websites and applications have also become much more complex with various portions of the end-user flow being impacted by an array of vendors. User login and authentication might be provided by an Identity and Access Management (IAM) provider, while DNS is managed by a DNS provider; a CDN might be involved in accelerating the entire web application; a user’s address might be validated by an API; credit card processing will be handled by one vendor while emailed reports are delivered by someone else.

In other words, many different vendors can take responsibility for different portions of the workflow. To complicate things further, each vendor might rely on others for specific portions of their services, creating an intertwined web of dependencies.

Not all observability necessarily enables oversight of all these vendors. Just because a specific service is being observed with a digital observability solution, it does not mean it is reachable, working, or the indicator measured by the SLA was not breached. This is due to the fact that services are being delivered by a highly complex architecture that involves multiple layers of providers, applications, infrastructure, and different networks.

A simple third-party service used by an application to validate and correct the user's mailing address, for instance, might rely on Oracle for the DNS, Akamai for the CDN and security protection, Amazon AWS for infrastructure, Amazon S3 for storing data, the code to be written by the service provider, etc. Each of these components tends to be highly distributed geographically with different transit ISPs and routing policies, and there might be hundreds of virtual machines on which the application runs.

In order to successfully manage the SLAs of multiple vendors, a company must:

- Ensure any SLA is sufficiently broad and robust.
- Capture observability data on critical business transactions.
- Observe the service using a system that actively tests the service and its outputs.
- Observe the service from outside the infrastructure hosting it in order to traverse the external network, the infrastructure and architecture of the service, and the application(s) handling the service requests.

Ensure your SLAs are broad and robust

In a recent [Gartner report](#) on SaaS SLAs, the research agency highlighted the fact that many SLAs are not “sufficiently broad or robust.” Gartner noted the key challenges of negotiating a robust SLA that needs to cover risk as broadly as possible, including failure to specify planned downtime exclusions and limitations, putting the customer at risk thereby of unplanned, and potentially unacceptable, outages at difficult times.

Gartner's recommendations included the need to:

- Utilize robust SLA vocabulary in SaaS contracts and incorporate applicable service-level targets;
- Negotiate financial penalties, which involve escalating tiered credits for missed targets;
- Negotiate the right to terminate if the agreed service levels are missed for three months within a 12-month period;
- Include the service-level reporting process in the contract itself, and ensure that the vendor is proactively alerting when reports are made available.



Observe your data in the right way

Sadly, most enterprises do not have a full understanding and/or visibility of the complexity of the relationships between their different service providers, and what this means to them. APM platforms offer visibility only once a major provider has a major outage and don't really help with monitoring SLAs.

One such event occurred in the fall of 2016 when the DNS provider Dyn, an Oracle company, had a major outage that took down the entire Internet. When the Dyn outage happened, it impacted not just its direct customers, but also the companies who relied on providers who relied on Dyn.

In other words, it was a domino effect. Most companies were not properly monitoring the full extent of their services, and only became aware of the problem after users complained that portions of their sites or web applications were not working properly.

These companies had no understanding, no visibility, and worst of all, no plan in place as to what to do when such an event took place. This widespread event gave birth to the **"multi-DNS" strategy** that most IT organizations have put in place since to ensure that DNS can no longer take them down.

Furthermore, an enterprise must ensure that its observability strategy observes the service level indicators involved in key transactions of their application or service; since each individual service involved in a transaction is capable of breaking it.

Observing transactions is more complicated than single page monitoring when determining which parameters should be used. Different inputs may yield different results based on application logic or the APIs used to pull information. Using the same inputs may not provide a sufficient level of insight.

For instance, it isn't feasible to test every permutation, but just checking that the page is up is not sufficient; however, validating that the entry of specific inputs results in specific outputs is very important.

Simply checking that the login page is accessible, for example, doesn't mean that the user is definitely able to login. You must enter a username and password, ensure the user is authenticated, and can reach a page with the correct data. For example, a certain kind of dashboard with a specific set of numbers and charts.

Catchpoint offers the ability to quickly create business workflows or end-user workflows as multi-step transactions through a Selenium-based Chrome extension.

The script recorder makes it easy to create transactions that can be used to detect any issues with key business processes. Logic can be inserted into scripts to choose different search terms, select valid travel dates, or by clicking on the second item in a list, let you dynamically cycle through a relevant list of terms, dates and/or numbers.

Additionally, Catchpoint offers not only observability of web transactions through a native Chrome browser that emulates that of end users, but also a myriad range of other services, from API transactions to DNS, email protocols to WebSocket, MQTT to network; even allowing for the building of customized monitoring for services unique to the individual enterprise.



Where you observe from matters

One of the biggest traps related to monitoring metrics for SLAs is where your observer probes are located. In the early days of SLA management, the service vendors figured out that if they observed their services from within their own datacenters, it was to their benefit.

However, buyers quickly learned there were mismatches between their vendor's SLA metrics and what the users of their companies were experiencing. In the early days of SLAs and Internet-based services, there were frequently cases in which the vendor's system operator would say everything was green (i.e., operating fine) while the technical support lines were ringing non-stop with furious customers who were experiencing outages.

Observing data 24/7 from a probe deployed on the same datacenter as the service might work for internal SLAs where the customers of the service are within the datacenter. However, this won't work for services consumed externally over the Internet.

A service on the Internet is more than just code on many containers on many physical servers. There are other components of the service architecture that are in the realm of control of the provider, which can cause serious performance issues, outages, and/or unreliable service transactions.

The key components that are always present are a datacenter's geographical location, transit providers used in the provider's datacenter, and routers, load balancers, and firewalls.

Lastly, the service provider itself will inevitably rely on third party services in the architecture mesh, which might be unknown to the service consumer, such as CDN, site acceleration, DDoS protection, page optimization, image optimization, cloud computing services (Lambda Functions, S3 Storage, etc.), and other services issued by the provider in other datacenters and so on.

The service provider ultimately controls which vendors are used and applies its own strategies to managing these services and enforcing its own SLAs. Considering that a modern service relies on so many components external to the datacenter, it clearly shows that solely measuring from within is not sufficiently reliable for SLA management.

At the same time, observing from each end user through **Real User Observability** or similar means will not work in most cases. At the edge of the Internet where the end users are, there are many components outside the control of the service provider and even that of the buyer.

Today an end user (customer or employee) might use a laptop and connect to a coffee shop WiFi, going through a Consumer ISP and an Internet proxy service. Any one of these components can introduce an outage or poor performance from an end-user perspective, but the provider cannot be held accountable for their failure. At the same time, Real User Observability will only have data available if the user is able to successfully access the service. Therefore, you won't have a constant stream of data available 24/7 to properly ascertain if an SLA has been breached. Moreover, when an outage occurs, you will have no data at all.

Since we cannot observe from within the datacenter or from the end user, the only place left is somewhere in between the provider's datacenter and the end user where the provider can actually be held accountable.



Observing from a single location isn't effective given the varied geographical ISP distribution.

You need to measure from a range of vantage points that match where your users are geographically located to ensure visibility of network issues that are in the hands of the provider.

The wider number of vantage points, the easier it is to see if issues are regional or global, and to have enough coverage for the distribution of the service, which could be in multiple datacenters and rely on multiple transit providers.

At the same time, having geography diversification is not enough. In different geographies, there are different transit providers that are key to the delivery of Internet data. Any service provider needs to ensure their network is accessible from all key backbone locations where communications will traverse through.

One of the latest observability solutions offered by digital observability vendors is active monitoring from the cloud. Such solutions offer a cost-effective way to monitor issues caused by your code. It is okay to measure from the cloud if that is where your application is hosted, but that is still only one piece of the puzzle.

APM and observability platforms, offering active monitoring solely from the cloud, will continue to fall short when it comes to SLA enforcement because:

- Cloud datacenters are frequently in geographies not close to large populated areas, because cloud providers will always find cost effective locations.
- Cloud providers buy transit from specific providers and thus have limited ISP diversification to catch issues that originate from those providers.
- Most importantly, the majority of your service providers are either hosting their services in the cloud or moving there. As a result, you are observing SLAs from within the infrastructure and services of the provider. As we have discussed earlier, this doesn't provide sufficiently reliable data.

The biggest SLA trap

The biggest trap in SLA management happens when companies negotiate each vendor SLA separately, and the negotiation is too often not conducted by those who are actually in charge of the design and architecture of the service. This results in the company not looking at what the SLA of a given vendor means to the overall service SLA.

To illustrate the problem, let's look in more detail at an example we briefly touched on earlier: a company has built a web-based application for selling data intelligence; its customers can purchase datasets as needed.

The service relies on:

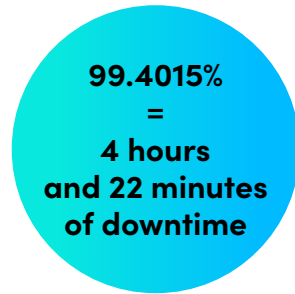
- DNS, which is managed by an external DNS provider.
- Web application is accelerated by a CDN provider.
- User login and authentication is provided by an Identity and Authentication Management (IAM) provider.
- Credit card processing is handled by a payment vendor (customers pay on a per transaction basis).
- The application has been deployed on a public cloud provider.

The company has a practice of setting all their availability **SLAs to 99.9%**, or 43m 49s of downtime per month. They decide to add a new vendor for DDoS protection, which sits between the CDN and the cloud provider; all transactions are shielded by the DDoS protection service. The team assumes they can negotiate an SLA of 99.9% as they have done previously.

However, they do not take into account that they themselves have internal SLAs with customers and partners. They mistakenly assume that they are covered, and will not have availability below 99.9% themselves. In reality however, the SLA of their application would be **breached at 99.9%** if more than one vendor breaches.

Their highest risk is if all of them breach, which would mean:

99.9
x 99.9
x 99.9
x 99.9
x 99.9 x 99.9



To avoid breaching their own SLAs with customers, or to reduce downtime for their users, a company must figure out how to reduce the risks beyond what the SLAs with each vendor allows them.

The need for a reliability strategy

SLA management is not just about holding vendors accountable. It is also about your IT organization ensuring reliable services independent of any vendor failure. This responsibility increasingly lies in the hands of Site Reliability Engineers (SREs), a discipline first introduced by Google, which has become significantly more prominent among sophisticated IT organizations in recent years.

For the last several years, we have been surveying SREs to understand more about this emerging role and any outages, incidents, and post-incident stress they encounter in their work.

In our [2019 SRE Report](#), **49% of respondents** stated they had been involved in incident resolution within the last week alone. When vendor failure happens, your organization should have in place a “vendor reliability strategy” by either implementing a multi-vendor reliability strategy or, depending on the service, a backup reliability vendor strategy.

Either way, you must **have an observability strategy** that supports your SLA strategy and a reliability strategy that involves real-time monitoring and alerting of the vendor’s service and yours.

How to implement a multi-vendor reliability strategy

In a multi-vendor reliability strategy, your organization ensures it has at least two vendors in an active-active mode all the time for the same service. This strategy allows IT organizations to architect their applications and services in such a way that the redundant vendors make up for any failures of their counterpart.

There are four key services in which this strategy has been successfully implemented:

- DNS
- CDN
- ISP peering
- Cloud compute

There are several challenges with implementing such a strategy, however. Your team needs to take these into account before going forward with this kind of approach:

- **Costs associated with dual vendors** – you will lose the economies of scale gained by going with a single vendor and incur double the usage rates.
- **Complexities involved in implementing and managing such a strategy** – it often requires building internal tools, processes, training and/or purchasing another vendor to manage the two services.
- **Failure to implement a monitoring strategy that supports this strategy** by assuming that implementing the multi-vendor strategy removes all risks; sadly, it does not.
- **It cannot be applied to all services.** For example, employee email might be using Microsoft Office 365, but there is no way as of today to have an active second provider, such as Google Workplace.

Why to consider a backup reliability strategy

A backup reliability strategy, by contrast, means that you have a backup plan in place for each key service which is outsourced to external vendors. Such a strategy is simpler and more cost effective, but you will still likely experience some impact from outages.

The backup strategy might be used in any one of these cases:

- Backup is to internal service(s). For instance, for a company portal with a single email provider for sending email alerts and reports, they can deploy mail servers in their infrastructure that they would use if the primary vendor goes down.
- Backup would be a second vendor, to be used in the case of an emergency. This will involve a lower spend but higher overages. In this case, you would have to make a small purchase for the second vendor in case of emergency. Your team puts the second vendor live only in instances of outages that impact the primary vendor. The ongoing cost of this approach can be significantly lower. However, it can still result in overages. CDNs and credit card processing are ideal services to do this for.
- Backup simply involves dropping the vendor's service. There are cases in which you can do this and everything will still function normally (or only non-key functions of the user flow will no longer be available). Good candidates for this are CDN, page optimization, tag management services, DDoS protection, or any service that won't impact actual business transactions if dropped or turned off.

What is a rock-solid observability strategy

In today's cloud era, no company can implement a successful digital service without a rock-solid observability strategy and an SLA enforcement policy to support it. It would be like building and running a business without having a dashboard of key performance indicators on how the business is operating.

Your monitoring strategy must not only support the ability to observe your vendors' SLAs, but also support your team in the implementation and management of any reliability strategies in place for digital services.

Your observability strategy must include the following capabilities:

- Actively monitor 24/7 (at a frequency of a minute or faster) an external DNS provider. Observe within your environment and outside it.
- Observe from all the key geographies where your users are located.
- Observe from key transit ISPs in those geographies.
- Observe key components of any digital service: DNS, network connectivity, HTTP, web transactions, email, Websocket, MQTT, etc.
- Observe key services that are handled by outside vendors: DNS, CDN, cloud, APIs, SaaS, email, etc.
- Provide real-time data and alerts based on captured data.
- Utilize real time APIs and integrations with other tools used in multi-vendor strategies.



Conclusion

The rise of digital services and the digital transformation movement in the enterprise typically relies on the use of cloud-based services provided by external vendors.

This has increased the risk of poor user experience and a negative impact on business, which in turn has pushed organizations into the robust use of Service Legal Agreements.

However, many companies today continue to lack proper Service Level Management practices. As a result, they are hit by adverse consequences when a breach occurs because they don't possess the means to hold their providers to account. The unfortunate outcome is that they are losing valuable customers.

Successful IT organizations have also reacted to potential vendor challenges by implementing multi-vendor or backup reliability strategies.

All these approaches, however, can only be successful alongside a broad and robust digital service observability strategy, which allows companies to observe their internal and external vendors to both hold them accountable and mitigate any outages in real time.

Depend on an independent arbiter to monitor and manage your SLAs:

www.catchpoint.com/sla-management

catchpoint.

Catchpoint is the enterprise-proven Digital Experience Observability leader. By providing unparalleled visibility and insight, we empower teams to confidently own the end-user experience. Learn more at www.catchpoint.com.

© 2022 Catchpoint Systems, Inc. All rights reserved. 210007-v1

