Benchmark CLOUDFLARE HashiCorp virtana

Immutable Clouds

MASTERING THE ART OF SITE RELIABILITY ENGINEERING (SRE)

Applying SRE principles to achieving the impossible – *100% availability*

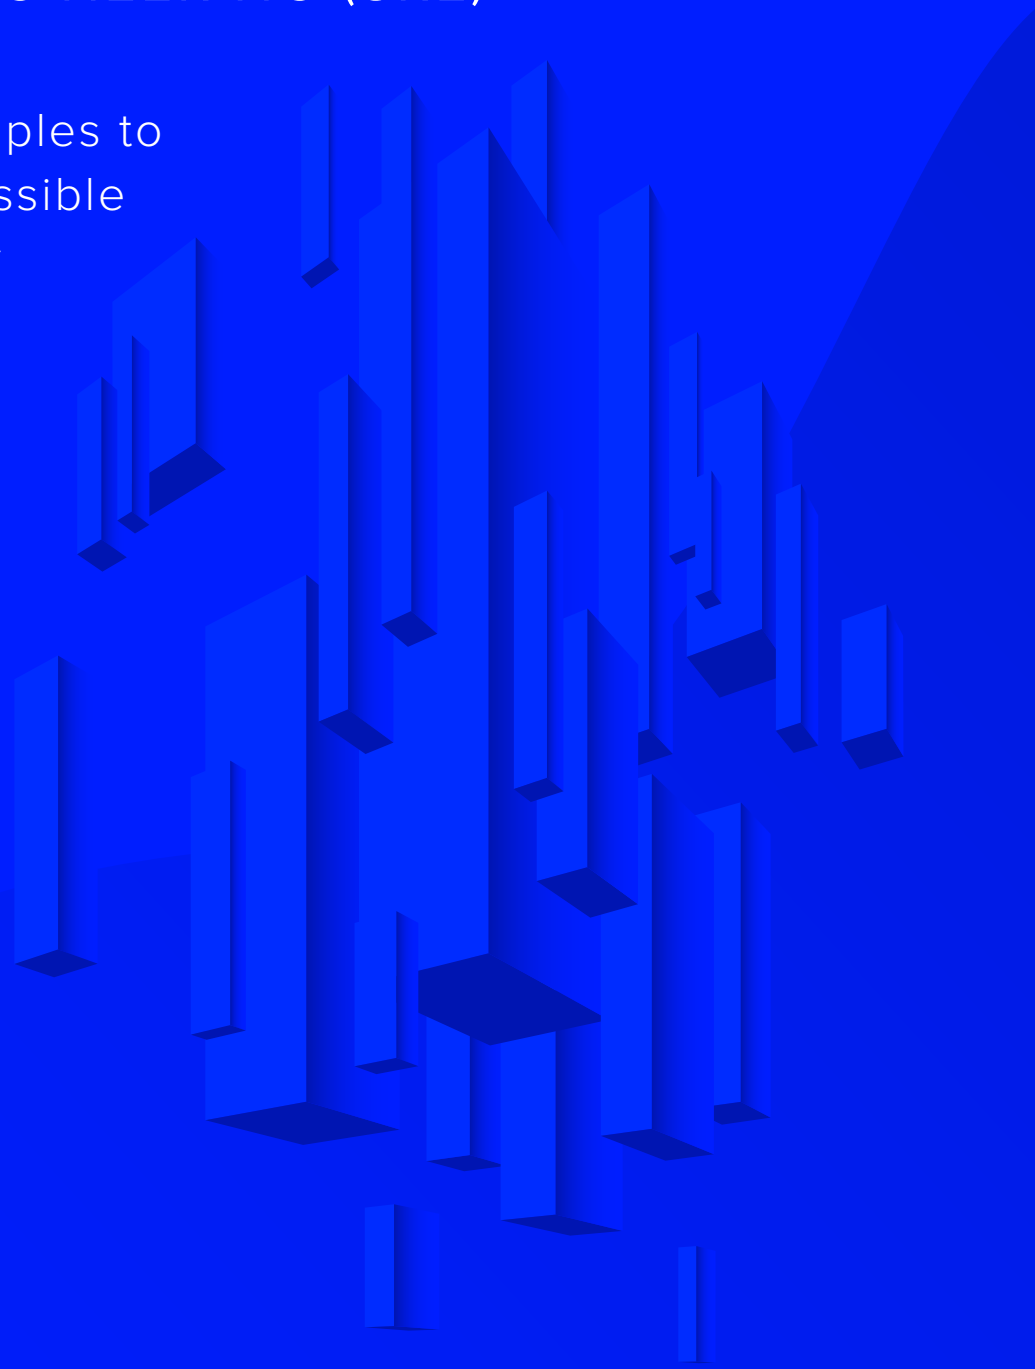

SHLOMO BIELAK
MATT COLE
ANDREW GNOINSKI
TOM BRIGHTBILL
JIM SULLIVAN

Introduction

With the adoption of cloud-native architectures, and platform as a service (PaaS) we have entered an era of immutable infrastructure and microservices-based applications. This has necessitated the need for chaos engineering practices and improved testing. Resilience to malicious attacks is now considered standard practice and keeping the same mutable machine running forever is not advantageous. The idea of live migration technology and resource utilization scheduling to prevent a failure from occurring is now seen as a hinderance to improving application resilience. This shift and advancement have provided Site Reliability Engineers (SRE) with a platform to implement an 'immutable cloud architecture.' The intent being to achieve the highest availability, for the longest timeframe possible, or a mean-time-to-restore (MTTR) that is not perceivable by the user. A single public cloud's service level agreement (SLA) is no longer enough for the SRE elite – 99.95%. One of the top cloud providers has stated that 100% availability is impossible.[1] In this white paper, we will explore what it would take to achieve the impossible feat of 100% availability across multiple cloud providers (private and/or public), and do so within a timeframe (months or years) using data and an updated SRE operating model. We believe this is our strength – providing our clients with a simplified technical sprawl and 50% reduction in operating costs.

[1] https://landing.google.com/sre/sre-book/chapters/service-level-objectives/

TABLE OF CONTENTS

SRE Operating Model

Architecture

Stateless micro-services are one part of the architecture equation in approaching 100% availability. The rapid adoption of containers is replacing expensive, mutable infrastructure. Cloud infrastructure is helping to accelerate this even faster.[2] Continuous delivery pipelines that use popular deployment strategies such as Blue-Green, Canary or Rolling have enabled a software delivery model to rapidly, securely and safely deploy applications. The SRE's operating model must extend beyond the application itself and apply the same principles to infrastructure automation pipelines, treating the entire cloud infrastructure as immutable – a fundamental prerequisite to improving your SLA.

> THE RAPID ADOPTION OF CONTAINERS IS REPLACING EXPENSIVE, MUTABLE INFRA.

Applying the standard gitOps model to infrastructure as code (IaC) provides the same level of governance and agility that is already proven for continuous delivery of applications. Branching, versioning and peer reviewing are standard best practices in application development. GitOps encompasses triggering infrastructure pipelines from git commit and initiation Software Development lifecycle (SDLC) practices. Similar to the software delivery model, when we treat infrastructure like code, we benefit from faster delivery, greater security and consistency. Applying an infrastructure as code methodology enables cloud portability, enhances availability, and avoids concentrated risk by reducing dependency on any single cloud provider. Cloud portability is a prerequisite to treating cloud as a commodity for technology consumption. Evaluating multiple clouds for your workloads allows for even greater flexibility and reliability, and inches even closer to that elusive 100% SLA.

ADOPTION OF ENGINEERING CLOUD TRENDS

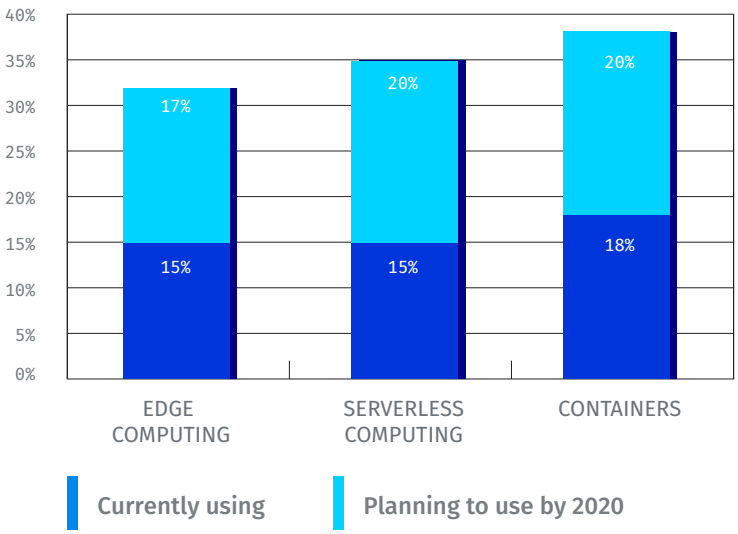

[2] **Figure 1:** Spiceworks – Public Cloud Trends in 2019 and Beyond

INFRASTRUCTURE PIPELINE

Figure 2

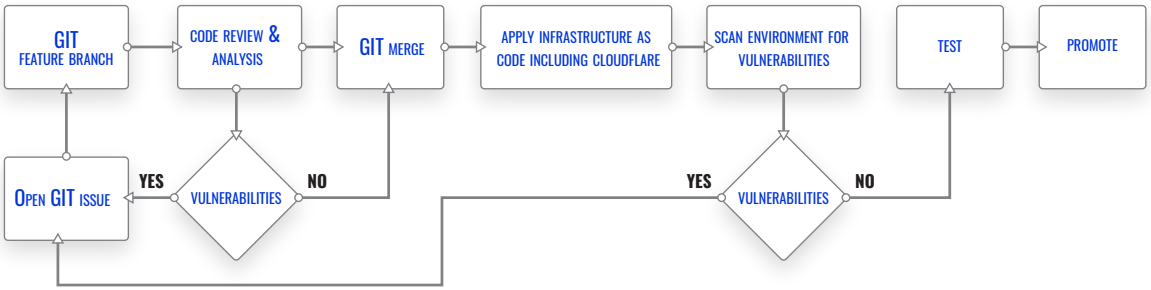

APPLICATION PIPELINE

Figure 3

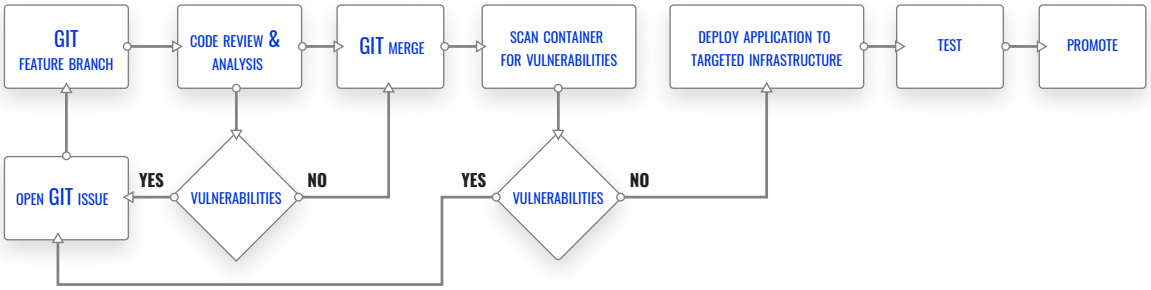

> "INTEGRATED AS PART OF THE INFRA PIPELINE, CLOUDFLARE GIVES US TRUE GLOBAL LOAD BALANCING."

Now that we have a reliable method to span infrastructure across multiple clouds, a mechanism is required to join them together. This approach is to support any workload running on any infrastructure in the cloud, compute or Kubernetes, as long as it has an ingress point. Overlaying a technology like Cloudflare allows for seamless distribution of traffic and workloads. Cloudflare gives us true global load balancing and is integrated as part of the infrastructure pipeline. Once you are able to balance connection of your workloads across multiple clouds, the user's experience is no longer tied to the uptime of one particular cloud provider, and your resilience increases. This will also reduce the leading cause for businesses to halt their use of cloud vendors; unreliable service.[3]

REDUNDANT MULTICLOUD ARCHITECTURE

Figure 4

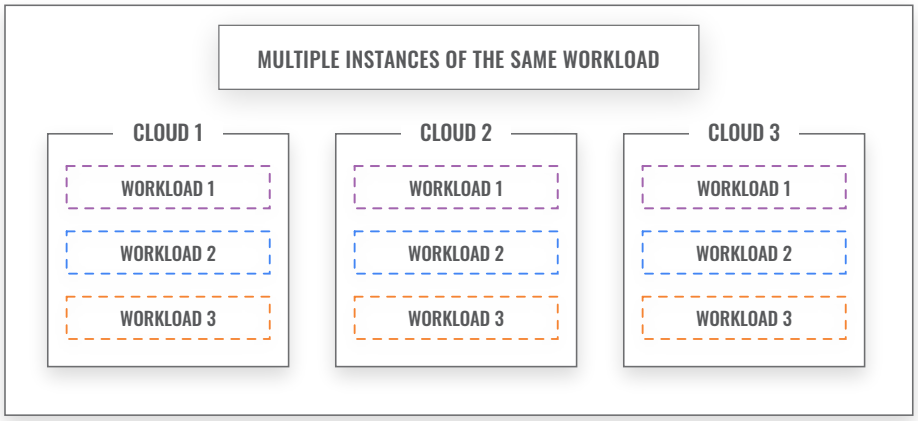

Top Drivers Leading Businesses to Stop Purchasing from Cloud Vendors

[3] **Figure 5:** Spiceworks – Public Cloud Trends in 2019 and Beyond

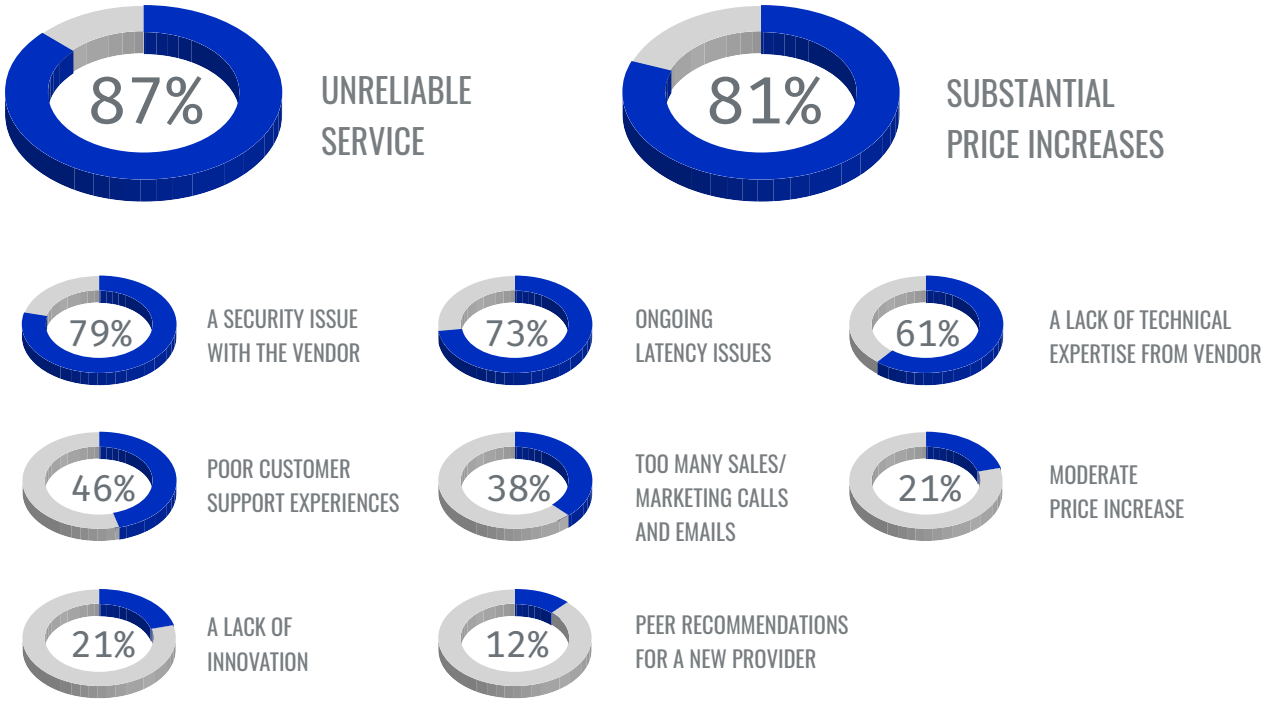

> YOU CAN CONTINUE TO QUICKLY, SAFELY AND SECURELY PROVIDE YOUR SERVICES TO YOUR CUSTOMERS.

Connecting your customers and users to your new architecture requires an operating layer above the cloud providers. The addition of another physical hop into the user experience to deliver services has been optimized by the CDN (content delivery network) offering and their SLA agreement. Their presence is vast and reliable. For example, Cloudflare's network has local Points of Presence in over 200 cities which serves a customer within 100ms. Furthermore, should an outage ever affect a single PoP in their network, the request would just fail over to the nearest location. The vast footprint of the CDN's network means you can continue to quickly, safely and securely provide your services to your customers.

ARCHITECTURAL EXAMPLE

Figure 6

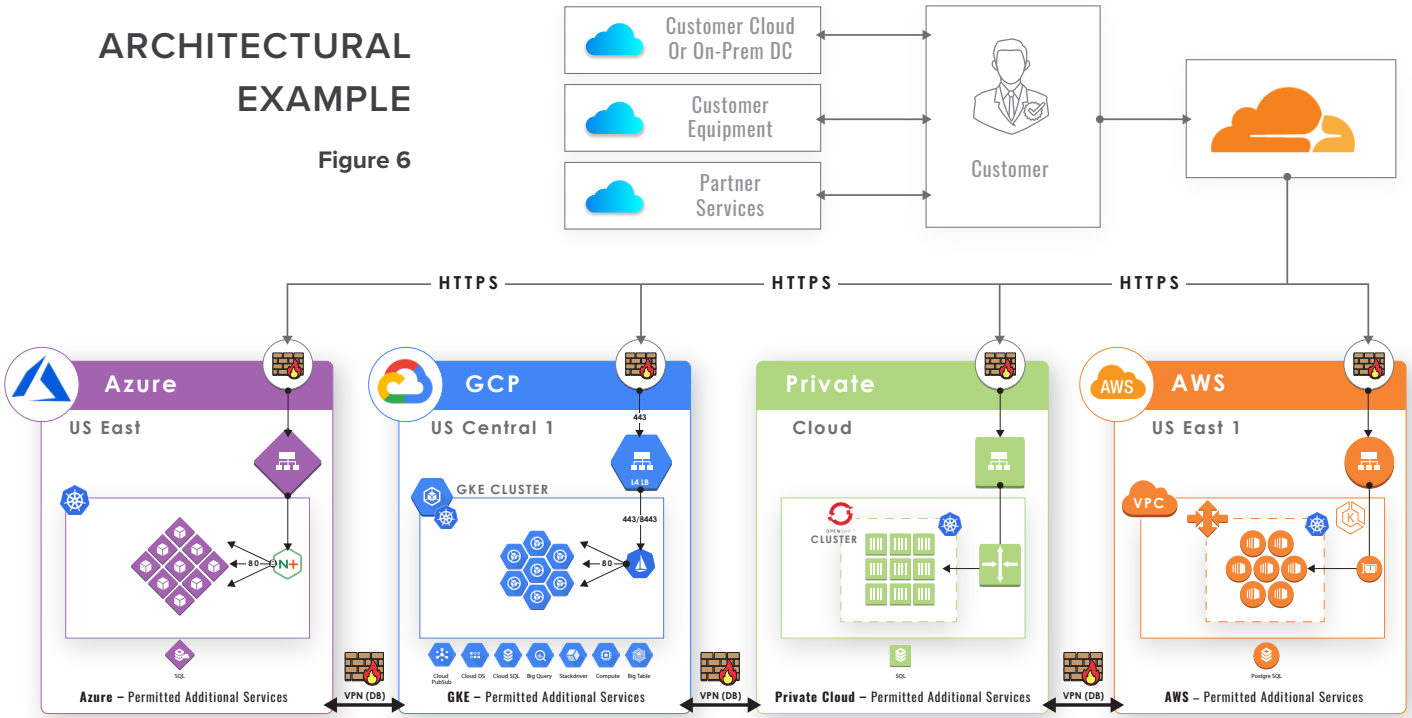

> ...CLOUDFLARE WOULD AUTOMATICALLY COORDINATE THE REDISTRIBUTION OF TRAFFIC ACROSS THE REMAINING AVAILABLE CLOUDS

The architectural example constructs the cloud infrastructure with a common declarative language, in this case HashiCorp Configuration Language (HCL) through Terraform. As part of that infrastructure, if a cloud outage or degradation occurs Cloudflare would automatically coordinate the redistribution of traffic across the remaining available clouds. Both Terraform and a CDN like Cloudflare support cloud agnostic sprawl, where infrastructure and services are distributed across multiple public and private cloud providers.

For hybrid cloud computing to work, there needs to be consistent workflows that can be reused at scale across multiple cloud providers. Therefore, these resources require:

- Consistent instruction sets for provisioning
- Identity and authentication management for security and network connections
- Privileges and rights so they can be deployed and run

> A COMMON CLOUD OPERATING MODEL IS AN INEVITABLE SHIFT FOR ENTERPRISES AIMING TO MAXIMIZE THEIR DIGITAL TRANSFORMATION EFFORTS

The essential implication of the transition to any cloud is the shift from “static” infrastructure to “dynamic” infrastructure: from a focus on configuration and management of a static fleet of IT resources to a focus on provisioning, securing, connecting, and running dynamic resources on demand. A multi-cloud operating model is an inevitable shift for enterprises aiming to maximize their digital transformation efforts. HashiCorp provides solutions for each layer of the cloud to enable enterprises to make this shift to the cloud operating model. Terraform focuses on the provisioning part of the HashiCorp Cloud Operating Model, enabling SREs to provision resources across this hybrid environment. [4]

But would using three clouds cost us more?

Let us explore the second leading cause for businesses to halt their use of cloud vendors; increasing costs. Having workloads in multiple availability zones is a common best practice. Our recommended solution goes beyond a singular cloud approach and allows for the distribution of workloads between multiple cloud providers in a single availability zone, with each having an equal footprint of capacity. Uptime and the features available are no longer tied to one provider but can be balanced across multiple cloud providers.

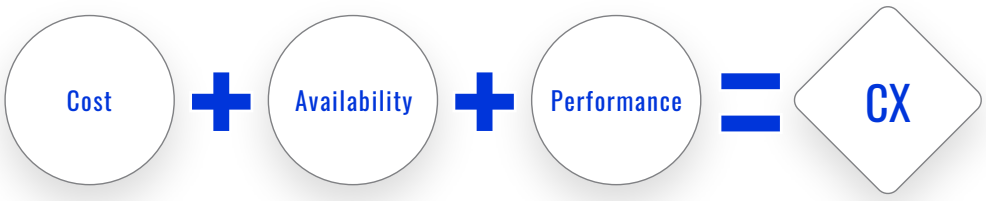

Figure 7

[4] https://www.hashicorp.com/cloud-operating-model

> THE NEED FOR CHANGE IS TO SUPPORT, IMPROVE, OR MAINTAIN THE CUSTOMER EXPERIENCE AND OPERATING COSTS

How do we know which technology or service/provider will be better for availability

As soon as your infrastructure and software teams have adopted and incorporated this architecture into their release model, an application can be delivered to a cloud of your choice based on an equation that provides the best possible customer experience with the highest business value. It is a balance between OPEX/CAPEX and required availability levels that your users and customers expect.

As a member of the partner advisory board for Virtana, Benchmark has been providing guidance to Virtana to help address a number of site reliability engineering use cases. An SRE requires data points to evaluate how a change will impact the customer experience. The SRE monitoring determines how a proposed change will affect the Service-Level Objectives (SLO) or SLAs that the organizations have with their customer. The most common sources of those changes are application releases; however, we also need to take into account infrastructure change and new technology adoption.

Change is inevitable and the need to be prepared to support, improve or maintain the customer experience while balancing operating costs is expected to run a growing business.

An effective SRE monitoring platform should provide all the necessary data points to make an informed decision for making a change to an application workload and the supporting infrastructure. The purpose of all applications is to serve your customers, requiring them to be available when your customers want to use the application, often 24x7x365. Therefore, we need to look comprehensively at the availability of an application, the availability of its infrastructure, as well as the overall cost. Based on historical data one can derive the MTTR for a particular service or application and evaluate if there is a tolerance for more change that would have business value (referred to as managing error budgets).

To deliver the best outcome, the decision about where to deploy and when to make a change needs to be based on SLAs, costs, and if available, historical data.

Virtana provides these data points to help streamline the CI/CD pipeline (API call) and simplify the decision **for where a new workload should go**, or if a change to an existing workload should be executed. Let's analyze both the scenarios.

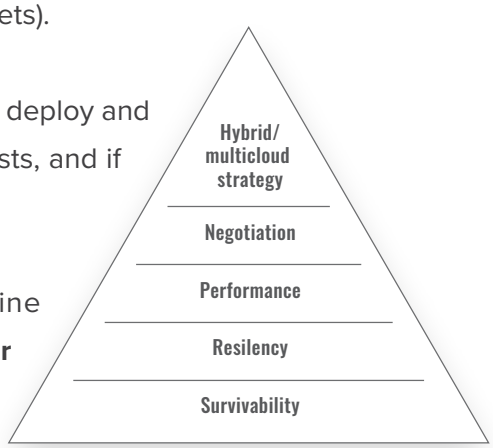

THE PYRAMID OF PORTABILITY AND MULTICLOUD NEEDS

Figure 8 © 2020 Gartner Inc

New deployment

WHERE SHOULD IT BE DEPLOYED?

> THE PURPOSE OF ALL APPLICATIONS IS TO BE AVAILABLE TO SERVE YOUR CUSTOMERS.

Today, clients decide on new workload deployments based on the features available at the public cloud providers and the discounts being offered by each cloud provider. Clients choose a single cloud provider for their deployment, without taking into account the right data and the consideration of application portability. We believe that every organization is going to pick a single cloud to start with; however, we recommend using data to drive the decision for where to deploy first and next. Knowing the expected service availability, historical performance of the application, optimal instance size, workload characteristics and expectations for growth will ensure you choose the most cost-effective cloud deployment for the target workload. This allows for an informed cost model to prevent an exit from the cloud while supporting a simplified multi-cloud adoption model.

Change to an existing workload

Delivering changes to multiple cloud providers can increase the layers of risk as there are vast data points that need analysis. We consider a change to be a disruption to the availability of the application. Changes occur often; however, it is not always a data driven decision. Hierarchy and change management policies ensure many eyes validate whether a change should happen. Data is gathered from a variety of tools, and combined with personal experience, to understand historical performance, how it is currently performing, how it is expected to perform in the future and how many outages or latency incidents have been experienced. A better approach would be to support SRE's and DevOps engineers with a platform that has all the relevant and necessary data to make the change decisions, as well as identify the available error budget for a workload, the number of recent changes and, if applicable, the cost savings that can be realized by changing the cloud instance size or cloud providers altogether. If an organization adopts multiple cloud providers, this data becomes more critical to sustaining workload availability.

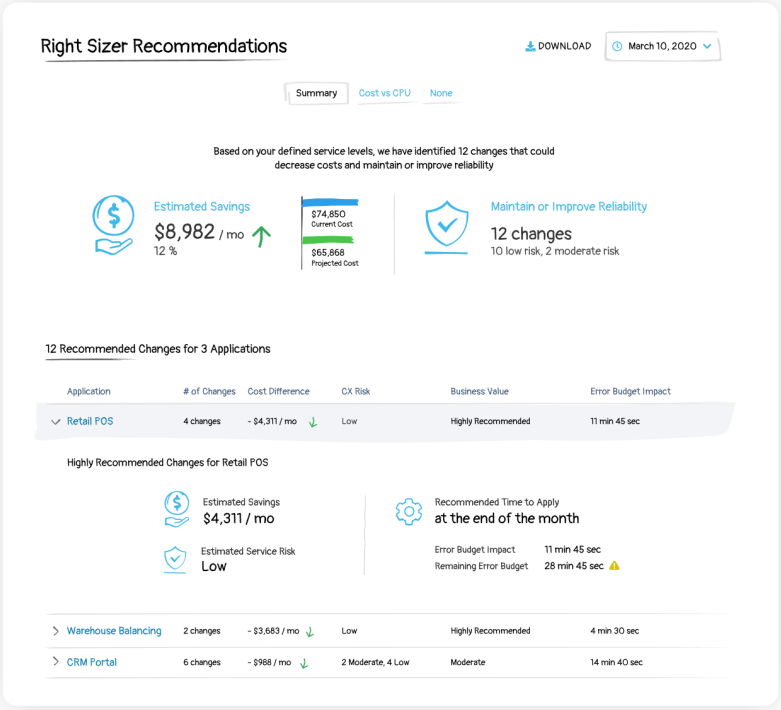

Figure 9

Better decisions can be made faster with the ability and choice to move applications between on-premise or multiple clouds based on the optimal deployment. Virtana provides clients with the availability and cost data needed for business case justification and answering the questions of why and where to deploy as well as when the change will have the least impact. This enables a shift in buying power from the cloud providers to you, the cloud provider's client.

You get metrics to decide where we should serve our customer from. Are there benefits to running in a cloud agnostic model, and does it really improve availability? The answer to this question is more complex as we must mature our orchestration layer to sprawl technologies with ease.

2015-2020 SLA Impacting Events on Cloud Core Infrastructure

(TOP 3 PROVIDERS) – DATA ANALYSIS

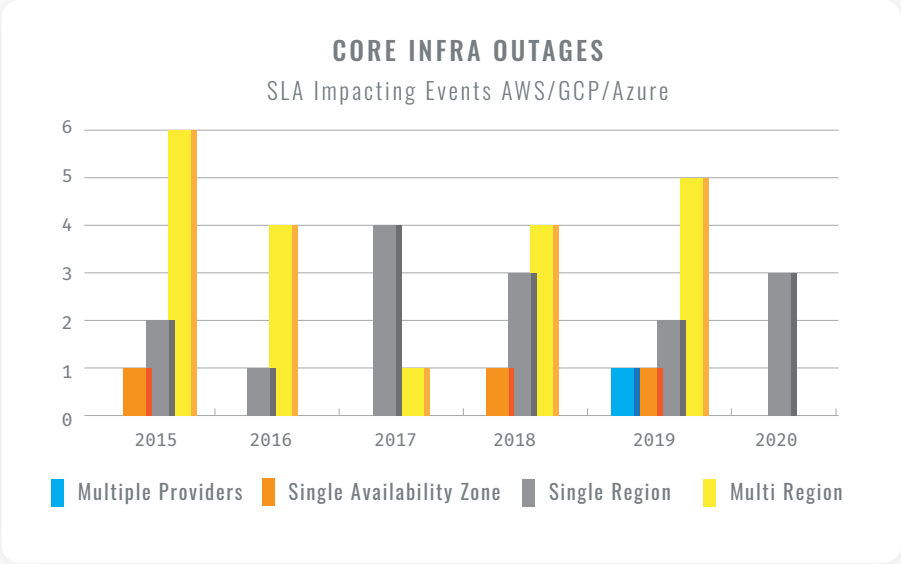

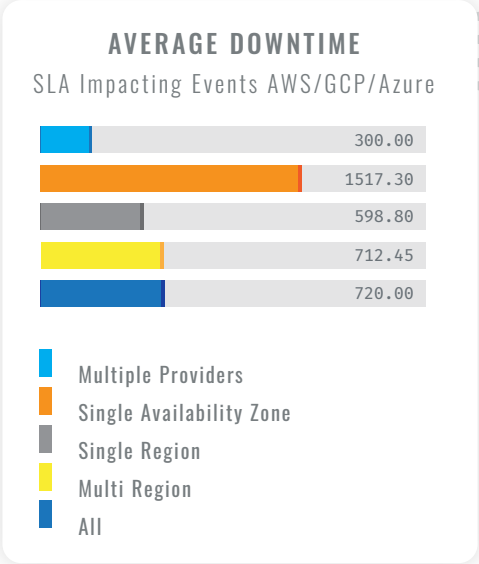

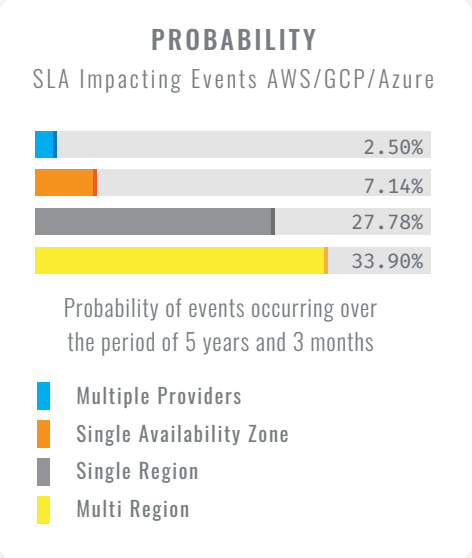

Figure 10

Let's explore the current state when using a single cloud provider. The most common model to sustain availability is to leverage multiple availability zones (i.e. three) and span an application across multiple regions. Most experts believe this provides the same coverage as having highly disparate (company and/or distance) data centers, which is perceived as the bellwether for business continuity. In terms of contract commitments and guaranteed SLAs, the availability uptime from cloud providers does not increase by adding regions or availability zones. Let's dive into the data points to validate the premise and, at the same time show where outages affect the top three cloud providers simultaneously.[5] We will analyze core infrastructure outages that caused downtime from 2015 to March 30, 2020.

[5] Thousandeyes, DownDetector, CloudHarmony, CRN, cloud provider status pages.

Data Analysis

> AN IMMUTABLE CLOUD ARCHITECTURE IS A REQUIREMENT FOR CRITICAL SERVICES AND THE DATA ANALYSIS PRESENTED SUPPORTS IT

Cloud Service provider outages can occur outside of guaranteed SLAs, impacting not only an availability zone, but an entire region or globally shared infrastructure. An immutable cloud architecture is a requirement for critical services and the data analysis presented in this white paper supports this position. The data shows single-AZ events (storm and power) are rare, being 8.6x less likely than regional/multi-region events (user error/updates). Furthermore, the suggested model does not require a regional distribution model to add additional resilience, **as no single region event or multi-regional outage at one provider affect other providers. The data summary shows resiliency designed within a cloud provider using multiple AZs and regions is more costly and less available than using multiple providers without those layers.** This would deliver significant OPEX savings, as the resilience to regional outages would be the single availability zone in another provider. Telco failures or DNS related issues can cause multi provider outages but the data shows this is very rare. Most outages that affect core infrastructure are lengthy due to the size and complexity of the cloud providers and affects their ability to find the root cause, further validating the strategy to have multiple cloud providers.

ONE CLOUD VS. 'SUM' CLOUD

Figure 11

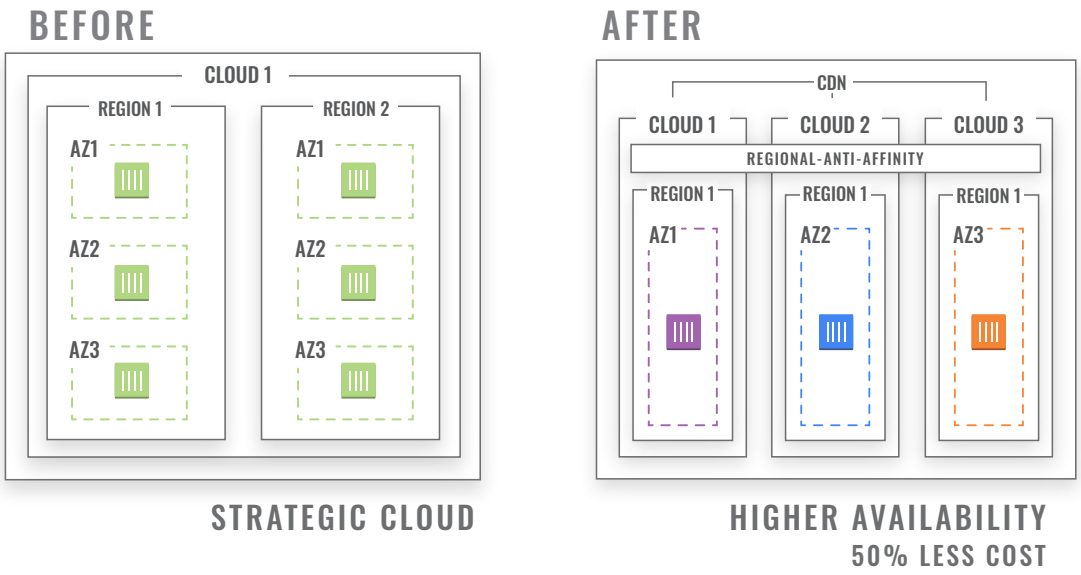

> **SINCE EACH CLOUD PROVIDER IS INDEPENDENT WE CAN REDISTRIBUTE THE LOST CAPACITY TO THE OTHER CLOUDS**

Business Continuity Plan

SUM CLOUD

The proposed architecture has multiple disaster recovery locations inherent in the design. However, we do not label any aspect of the architecture as recovery. The architectural design goal is to achieve 100% availability while not being impacted by a loss of an entire provider. The loss of a region is expected and managed by design, without additional cost. The loss of the data center or service is managed at the CDN's health awareness checks to prevent data going to a provider that is not servicing customers. The data supports the approach as a multiple provider outage (3 or more providers) is extremely rare and well within our SLA window or error budget.

The continuity plan is more in relation to the persistence of data and data replication requirements. There are methods to ensure the data continuity between providers using replication strategies to sync data sources between each cloud provider. This will allow for a DNS based database-connection-string to move with the CDN's health check. Similar to a Global Traffic Manager for databases. The design represents potential savings for a Business continuity plan (BCP) if there were additional or dedicated disaster recovery infrastructure. **Since each cloud provider is independent we can redistribute the lost capacity to the other clouds**. With a single cloud provider design this would not be possible due to increased risk using a common region.

BUSINESS CONTINUITY PLAN

Figure 12

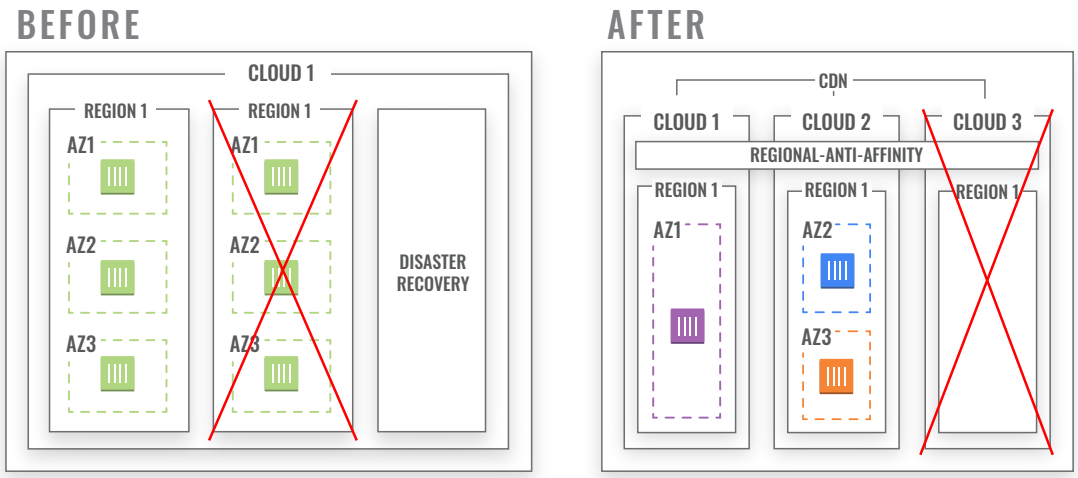

STRATEGIC CLOUD

HIGHER AVAILABILITY
DR BY DESIGN

Cementing the Sum Cloud Architecture

> A NEW HOME FOR CRITICAL SERVICES THAT OTHERWISE COULD NOT MIGRATE BEFORE

The data proves that a sum cloud architecture is the highest possible level of resilience with great cost efficiencies. By categorizing the outages, we found that there were multiple years that did not have an outage that affected the CDN or all three providers at once. The SLAs we publish yearly as accolades and battle scars could have been better, dare we say perfect, if we adjusted our SRE approach to incorporate cloud level infra-as-code. An elite SRE can stand up from scratch and deploy securely into all three cloud providers with micro-services in 30-45 minutes. Delivering redundant workloads into each cloud provider attaining the SLA a critical service requires, such as a core mobile banking application. Without the regional resiliency (proven ineffective) one could effectively attain a 50% OPEX savings using this model. The distribution of availability nodes between providers is cost neutral beyond telemetry aggregation/CDN. This architecture is a harness, through our DevOps battle scars, to provide a cost effective and resilient model for banks and other critical services offerings to advance beyond their current data centers. This call for change is for those that have already solved the application delivery model, SRE elites, as it requires those who have had their talent heroic scars heal over multiple years. Let's advance together and find a new home for critical services that otherwise could not migrate.

> ONE COULD EFFECTIVELY ATTAIN A 50% OPEX SAVINGS USING THIS MODEL

By abstracting the cloud provider layer, Cloudflare eliminates the risk and concern of a singular cloud outage, reduces MTTR, and gets even closer to 100% SLA. More importantly, the customer experience remains unaffected. If you leverage a robust SRE operating model that includes infrastructure pipelines, which follows our Benchmark motto of 'coding complex simply', you can attain an Open Cloud model without a large talent base. An SRE monitoring platform like Virtana would provide the data analytics on a real-time basis to target and deliver your services effectively. Integrating this with your CD pipeline to leverage any and all the cloud providers that benefit your customer experience. Using a delivery model for infrastructure and applications, the architecture for any cloud, and the data points for technology consumption provides high resilience, challenging Google's statement that 100% availability is impossible. Whether or not you need to achieve that milestone you will concurrently reduce risks and costs by abstracting the cloud layer.

Benchmark CLOUDFLARE HashiCorp virtana

Immutable Clouds

MASTERING THE ART OF SITE RELIABILITY ENGINEERING (SRE)

If you'd like more information or would like to set up a half-day complimentary workshop with your team, please contact

info@benchmarkcorp.com