Software Engineering in Industry

Nadine Abdelhalim, Yujiao Xiao, Luca Mehl
Word Count: 1991

I. Introduction

In this article we analyse the software architecture, cloud adoption strategy, and sustainability strategy of two retail SaaS companies: Shopify and Squarespace. We compare distinct approaches and assess advantages and drawbacks.

Shopify is a multi-tenant e-commerce platform that helps all sizes of merchants set up, manage, and scale webshops [1]. Squarespace provides tools and resources to create websites, set up online stores, and market products or services [2].

Chosen for their similar business contexts, as SaaS entities operating in the online retail sector, and comparable scale, Shopify and Squarespace differ notably in their architectural approaches [3], [4]: Shopify is a rarity among companies of its scale, deliberately embracing monolithic architecture [5], while Squarespace currently operates with a microservice architecture [6].

II. Software Architecture

Software architecture is the foundational design guiding complex systems. This section explores how Squarespace and Shopify strategically navigated their architectural transitions.

Shopify

A. Architectural Evolution - The Monolithic Approach

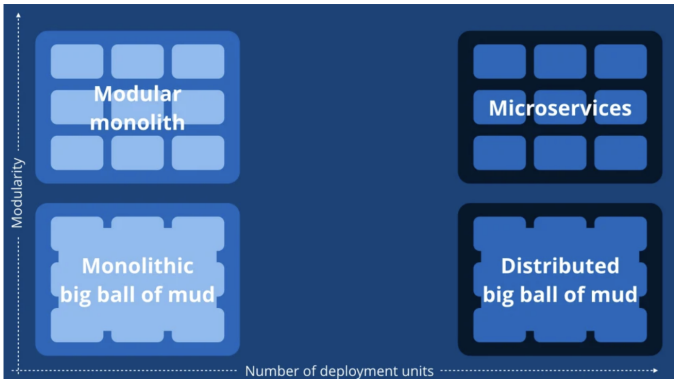

Fig. 1. Monolith vs Microservices by Simon Brown [5]

Initially, Shopify adopted a monolithic structure due to its simplicity and agility, the standard practice for many startups [7]. However, as the company expanded, so did its core monolith, accumulating over 2.8 million lines of Ruby code [8]. This growth posed challenges for the developers, especially in building and testing new features, with developer enjoyment declining YoY due to the high code coupling [9]. By 2016, having recognised these limitations, Shopify's development team sought to "componentise" their architecture to improve "velocity and enjoyment" [5], [10]. However, they recognised the challenge of a complete rewrite that transitioning to microservices would require [5], [10], and so, unlike the industry norm [11], they adopted a modular monolith. This approach improved their architecture by streamlining data access and simplifying infrastructure management, making potential future migration to microservices easier by creating well-defined boundaries within the codebase [8].

In fact, despite Shopify maintaining essential functionalities within the core monolith, the ongoing migration process has seen Shopify extract non-core components into standalone services [10]. The new architecture notably facilitated this migration, significantly easing the extraction process [8].

B. Why was this approach successful?

Shopify's journey provides invaluable lessons for organisations navigating similar growth trajectories. It highlights the significance of *tailoring strategies* to meet specific organisational needs and avoid the urge to blindly adopt trending tech solutions [7]. Additionally, understanding *developer behaviour* was critical for pinpointing genuine challenges; an architectural change will not be successful if it works "against the momentum of hundreds of developers adding features" [8], [10]. Their success is also attributed to an *incremental approach*, advocating gradual changes that align with ongoing feature development, facilitating organisational buy-in [10].

In our recent interview with a Shopify developer, they stated that the monolith could sometimes "make it slow to run tests", but the migration improved developer enjoyment and hastened overall development.

Shopify's experience suggests that companies should plan for *prolonged transitional phases*, as almost all large architectural changes end up *incomplete*; acknowledging and planning for this ensures adaptability and successful architectural shifts [8], [10]. Decide on an architecture that provides an "intermediary state that is most useful for your situation" [8].

Squarespace

A. Organisational Shifts and the Microservice architecture

Unlike Shopify's initial code-focused architectural shift, Squarespace's transformation began with an organisational restructuring, aligning teams around specific business areas rather than skill sets, mirroring Conway's law that correlates software systems with organisational communication structures [12]–[14].

Around 2014, Squarespace began its global expansion efforts and, as a result, initiated its move toward microservices [15]. They started with three simple microservices–location,

taxation, and currency–to develop the initial microservice framework.

Their pursuit of platform reliability prompted heavy investments in observability and alerting systems, as seen in Figure 2. [6], [16].

Fig. 2. Squarespace observability and alerting [6]

Diverse microservices resulted in more complex capacity planning and longer wait times for provisioning services. Conflicting database usage between services and the existing monolithic system required careful design to ensure service independence. Despite this, microservices improved developer productivity, accelerated functionality releases, and granted autonomy to individual teams [6].

B. Why was this approach successful?

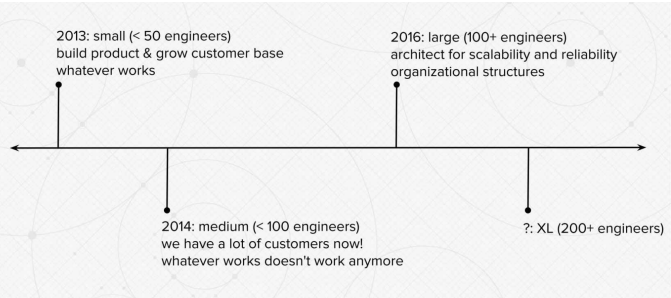

Fig. 3. Squarespace's growth timeline [6]

One of the key factors that aided Squarespace in this transition was its relatively *small organisation size*. At the onset of their shift, Squarespace counted roughly 100 engineers, a notable contrast to Shopify's workforce of over 5000 employees during their architectural transition in 2016 [6], [10]. Understanding developers and influencing their behaviour becomes notably more manageable at a smaller scale. The small size of Squarespace's organisation enabled its core services team to serve as an advisory body for all other developers, a task unmanageable within a larger organisation.

Furthermore, a key problem with microservices is the difficulty in tracking their state and failures; Squarespace's *meticulous focus on observability* allowed them to address emerging challenges rapidly, ensuring the adaptation process remained agile and responsive.

ARCHITECTURAL COMPARISON

Both companies initially adopted a monolithic architecture due to its swifter time-to-market advantage [7]; however, as they grew, mounting challenges led to a shift toward componentisation. Notably, Shopify and Squarespace pursued this architectural transition for different objectives: Shopify aimed to increase velocity and improve developer satisfaction, while Squarespace sought to improve scalability [9], [16].

While Shopify's modular monolith allows for easier management and avoids the complexities of a complete microservices transition, Squarespace's microservices-based architecture offers increased team autonomy and flexibility, albeit with challenges in standardisation [16], [17]. Both approaches aim to support scalability and maintainability as these platforms continue to evolve.

III. CLOUD ADOPTION

Cloud computing has significantly changed the tech landscape by providing a flexible and scalable option compared to the traditional on-site infrastructure [18]. This section discusses the cloud transformation of Shopify and Squarespace.

SHOPIFY

Today, Shopify predominantly hosts its infrastructure on Google Cloud Platform (GCP). However, this was not always the case [1], [8].

Initially, Shopify owned and maintained physical servers within their own data centre [19]. In 2015, they realised the importance of having a disaster recovery strategy and thus developed a duplicated data centre [19]. However, they found it was not cost-efficient [19]. Ultimately, they chose to move to the cloud, which enabled Shopify to meet data storage and privacy requirements, especially with the increasing regulations such as Safe Harbour [20]. Collaborating with Google allowed them to take advantage of its mature cloud infrastructure, solving their problems while gaining additional benefits.

By using cloud services, Shopify only needs to pay for the resources it utilises [19]. This allows Shopify to optimise its cost and provide a competitive tiered pricing model for users [21]. Moreover, the cost efficiency increases the scalability of their business.

Moreover, utilising modern tools like Spin and Isospin, and switching to an infrastructure based on Kubernetes enabled Shopify to simplify its development process by abstracting complex local development environments [21]. Thus, developers can efficiently develop new features without understanding the development environment [21]. In turn, this reduces time-to-market for new functionalities and allows Shopify to respond quickly to market demand.

SQUARESPACE

Until 2022, Squarespace maintained a self-hosted Hadoop ecosystem consisting of two autonomously managed clusters, designed in an active/passive model for geo-redundancy [22]. The ageing infrastructure and the demand for scalability led to mounting maintenance costs and scaling complexities [22].

Additionally, their adoption of a microservices architecture exacerbated these issues, significantly complicating data centre management. The microservices adoption demanded more frequent capacity planning for the SRE and TechOps Teams. Microservices introduced a wide array of services with diverse databases and programming languages. This caused delays in accessing staging environments, leading to substantial operational bottlenecks. Furthermore, the shift to microservices mandated diverse recovery strategies [6].

Squarespace, therefore, chose to migrate to the cloud. They adopted an iterative approach to ensure a smooth transition and only altered small pieces at a time; their final cloud architecture is shown in Figure 4. This saved time for manual validation at each migration phase. After migrating, Squarespace experienced an 87% drop in escalations. This allowed data teams to focus more on developing new features and improving their software instead of monitoring services [22].

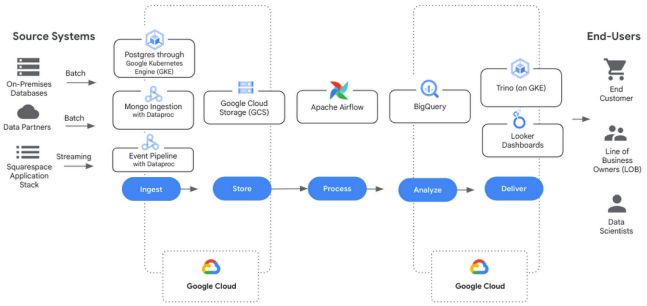

Fig. 4. Squarespace Cloud Architecture [22]

Comparison of Cloud Adoption Strategies and Factors Contributing to Success

Despite initially relying on on-premises infrastructure, both companies eventually migrated to the cloud, albeit for different reasons. Shopify pursued enhanced disaster recovery strategies, while Squarespace aimed to alleviate rising costs and growing complexities [6], [23].

Shopify's success stems from its recognition of the limitations within its expanding infrastructure. The company realised the inefficiency of operating its data centre due to escalating costs and expansion demands. Their decision-making process involved a meticulous cost-benefit analysis before embracing cloud migration [8].

On the other hand, Squarespace took a detailed and systematic approach to migration, creating comprehensive migration plans and transitioning small segments gradually [6]. This method facilitated a smoother transition, minimising risks and preserving operational continuity [18], [24].

Ultimately, their success in cloud adoption was rooted in their strategic alignment of business objectives with technological capabilities and their meticulous, step-by-step approach to migration.

IV. Sustainable Software

Anthropogenic climate change is driving global recognition of the necessity for sustainability in industry. Software engineering is no exception: a growing movement toward greener development, execution, and maintenance of software is emerging.

Shopify

The majority of Shopify's platform operates on GCP, which Shopify claims "runs on 100% renewable energy", which they imply absolves them from explaining the tremendous energy cost of hosting their platform [25]. However, this claim is highly misleading and far from a silver bullet.

Carbon Removal Market

Through their involvement in the Frontier Initiative, Shopify committed $11 million in 2022 to support the development of varied carbon removal technologies [26]. The initiative seeks to create a stable carbon credit market, catering to both suppliers seeking capital and buyers desiring simplicity and reliability [26].

However, while their carbon removal is promising, making changes to the business model may have a far greater impact in the online retail space; Shipping on Black Friday weekend alone accounts for 73000 tons of CO2e–equivalent to the annual emissions of over 15 *thousand* cars [25], [27].

Squarespace

Like Shopify, Squarespace defers environmental responsibility for their data centre usage to GCP. [14]

No-code sustainability

Beyond Squarespace's internal system, we must assess how Squarespace influences the energy efficiency of the websites it hosts. Squarespace's user-friendly platform allows non-coders to build sites effortlessly, hosting millions of websites on a unified software and hardware stack, thus benefiting from scale efficiencies and fostering dedicated code optimisation efforts. However, using low/no-code builders can result in heavier, data-laden websites, increasing carbon footprint [28]. In fact, Squarespace's own homepage generates 0.64g of CO2 per visit, ranking higher than 61% of global web pages in carbon emissions [29].

Sustainable Software Comparison

Not-so-Green Electricity

As seen in their respective ESG reports, Shopify and Squarespace have tried to outsource the responsibility of running their platform sustainably to their cloud provider, GCP. "Our platform runs entirely on Google Cloud, which means merchants' online stores are powered by renewable energy" [25]. However, this is *factually incorrect*. Like many of its competitor cloud providers, Google has adopted the misleading metric that it is buying enough renewable energy to 100% match its energy usage. In fact, they rely on legacy energy sources to fill the gaps when renewables are unavailable [30]. A more appropriate metric is the Carbon-Free Electricity percentage (CFE%), of which *only one region* on GCP has 100%.

Water Usage

Electricity is not the only resource consumed by data centres: water is crucial for server cooling, and a large data centre can use "anywhere between 1 million and 5 million gallons of water a day — as much as a town of 10,000 to 50,000 people." [31] For instance, the GCP 'us-west1' region in The Dalles, Oregon, consumes more than a quarter of the city's dwindling water supply [32]. In the UK, Google was investigated for excessive consumption during the 2022 summer drought. [33] By 2030, global freshwater demand will exceed supply by 40% [34], exacerbated by a warming climate, while demand for compute resources is expected to scale similarly [35]–putting software in competition with the global population for water supply.

Simple strategies to mitigate this carbon-intensive energy and water usage include spatial and temporal shifting [36]. However, a clear takeaway from our case studies is the trend for cloud providers to be opaque about environmental harm and companies to take them at their word. Thus, software developers should remain critical of sustainability claims and hold their own companies accountable where possible.

V. Conclusion

The analysis of Shopify and Squarespace's software strategies illustrates their nuanced approaches to scaling. Shopify's modular monolith sped up development, while Squarespace's microservices empowered teams and increased autonomy. Both shifted to the cloud for diverse reasons—Shopify for disaster recovery, Squarespace for cost and complexity management and reaping benefits in scalability and security. Yet, their environmental claims fall short; both rely on GCP, claiming renewable energy use, overlooking nuanced energy realities and water consumption issues.

References

[1] "How does shopify work? a guide for understanding this cloud-based webshop system," May 2023. [Online]. Available: https://www.magnalister.com/en/how-does-shopify-work/

[2] "Squarespace," December 2023. [Online]. Available: https://www.squarespace.com/

[3] Nov 2023. [Online]. Available: https://investors.squarespace.com/news-events-financials/investor-news/news-details/2023/Squarespace-Announces-Third-Quarter-2023-Financial-Results/default.aspx

[4] "Shopify announces first-quarter 2023 financial results; agrees to sell shopify logistics to flexport," May 2023. [Online]. Available: https://news.shopify.com/shopify-announces-first-quarter-2023-financial-results-agrees-to-sell-shopify-logistics-to-flexport

[5] K. Westeinde, "Deconstructing the monolith: Designing software that maximizes developer productivity," *Shopify. Feb*, 2019.

[6] F. Angulo, "Building microservices @squarespace," 2017, qCon Software Development Conference. [Online]. Available: https://www.infoq.com/presentations/sqarespace-microservices/

[7] D. Gravanis, G. Kakarontzas, and V. Gerogiannis, "You don't need a microservices architecture (yet): Monoliths may do the trick," in *Proceedings of the 2021 European Symposium on Software Engineering*, ser. ESSE '21. New York, NY, USA: Association for Computing Machinery, 2022, p. 39–44. [Online]. Available: https://doi.org/10.1145/3501774.3501780

[8] P. Müller, "Under deconstruction: The state of shopify's monolith," *Shopify. Sep*, 2020.

[9] S. Hanselman, "862: Engineering leadership and refactoring monoliths with kirsten westeinde," Oct. 2022. [Online]. Available: https://hanselminutes.com/862/engineering-leadership-and-refactoring-monoliths-with-kirsten-westeinde

[10] B. Liu, "Decomposing a massive rails monolith with kirsten westeinde, software development manager, shopify," Aug. 2021. [Online]. Available: https://about.sourcegraph.com/podcast/refactoring-shopify-codebase-kirsten-westeinde

[11] S. Newman, *Monolith to microservices: evolutionary patterns to transform your monolith*. O'Reilly Media, 2019.

[12] J. Herbsleb and R. Grinter, "Architectures, coordination, and distance: Conway's law and beyond," *IEEE Software*, vol. 16, no. 5, pp. 63–70, 1999.

[13] M. E. Conway, "How do committees invent," *Datamation*, vol. 14, no. 4, pp. 28–31, 1968.

[14] Squarespace, "Squarespace releases its inaugural environmental, social and governance report," Mar 2023. [Online]. Available: https://newsroom.squarespace.com/blog/squarespace-releases-its-inaugural-environmental-social-and-governance-report

[15] Apr 2014. [Online]. Available: https://www.squarespace.com/press-releases/2014/10/13/squarespace-raises-40m-in-series-b-funding-from-global-growth-investor-general-atlantic

[16] R. Kachouh, "The pillars of squarespace services," May 2017. [Online]. Available: https://engineering.squarespace.com/blog/2017/the-pillars-of-squarespace-services

[17] S. Esparrachiari, T. Reilly, and A. Rentz, "Tracking and controlling microservice dependencies: Dependency management is a crucial part of system and software design." *Queue*, vol. 16, no. 4, pp. 44–65, 2018.

[18] G. S. Sunilkumar Manvi, *Cloud Computing: Concepts and Technologies*. " O'Reilly Media, Inc.", 2021.

[19] S. Francis, "Shopify's move from the data centre to the cloud," June 2018. [Online]. Available: https://www.usenix.org/conference/srecon18asia/presentation/francis

[20] A. Philopoulos, "Squarespace releases its inaugural environmental, social and governance report," Apr 2020. [Online]. Available: https://www.shopify.com/uk/blog/shopify-experts

[21] D. Kelly, "The journey to cloud development: How shopify went all-in on spin," Jun 2022. [Online]. Available: https://shopify.engineering/shopifys-cloud-development-journey

[22] D. O'Connor, "Squarespace reduces number of escalations by 87 percent with analytics lakehouse on google cloud," April 2023. [Online]. Available: https://cloud.google.com/blog/products/data-analytics/squarespace-migrates-from-on-prem-hadoop-to-google-cloud

[23] D. Neufeld, "Shopify's infrastructure collaboration with google," Mar 2018. [Online]. Available: https://shopify.engineering/shopify-infrastructure-collaboration-with-google

[24] "Adoption of cloud computing as innovation in the organization." [Online]. Available: https://journals.sagepub.com/doi/full/10.1177/18479790221093992

[25] "Shopify 2022 climate report." [Online]. Available: https://www.shopify.com/climate-report

[26] "Frontier climate," December 2023. [Online]. Available: https://frontierclimate.com/

[27] "Greenhouse gas emissions from a typical passenger vehicle," Aug 2023. [Online]. Available: https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle

[28] M. Moran, "Squarespace statistics (how many websites use squarespace?)," April 2023. [Online]. Available: https://colorlib.com/wp/squarespace-statistics/

[29] "Website carbon results for: www.squarespace.com," November 2023. [Online]. Available: https://www.websitecarbon.com/website/squarespace-com/

[30] "Carbon free energy for google cloud regions," November 2023. [Online]. Available: https://cloud.google.com/sustainability/region-carbon

[31] S. Osaka, "A new front in the water wars: Your internet use - the washington post," Apr 2023. [Online]. Available: https://www.washingtonpost.com/climate-environment/2023/04/25/data-centers-drought-water-use/

[32] M. Rogoway, "Google's water use is soaring in the dalles, records show, with two more data centers to come," February 2023. [Online]. Available: https://www.oregonlive.com/silicon-forest/2022/12/googles-water-use-is-soaring-in-the-dalles-records-show-with-two-more-data-centers-to-come.html

[33] D. Robinson, "Thames water investigating datacenter water use in london," Aug 2022. [Online]. Available: https://www.theregister.com/2022/08/24/thames_water_datacenter_water

[34] F. Harvey, "Global fresh water demand will outstrip supply by 40

[35] D. T. Kari Beets, Glen Duncan, "2023 global data center outlook," April 2023. [Online]. Available: https://www.jll.co.uk/en/trends-and-insights/research/data-center-outlook

[36] “Carbon awareness.” [Online]. Available: https://learn.greensoftware.foundation/carbon-awareness/