

The Emergence of Supercloud

How SaaS platforms in the post-cloud hyperscaler era are turning the multicloud mess from a CIO's biggest problem into today's biggest opportunity

“Multicloud”. To a public cloud salesperson or executive, it’s practically profanity.

But to CIOs and technical leaders, it’s nearly as unloveable. The prospect of maintaining their on-prem data centers, migrating to one public cloud, *and* redundantly staffing for two or more clouds is mind numbing.

And let’s be honest: Most won’t achieve a viable multicloud strategy any time soon. There’s a big enough burden with the cost, complexity, and staffing challenges of just getting two clouds (private and one public) to interact.

This is despite the strong and understandable reasons to get to multicloud: Avoiding vendor lock-in, access to best-of-breed services regardless of vendor, and the flexibility to successfully interact with business partners, M&A-acquired IT stacks, etc. *If ever there were a need for an innovative alternative from the market, this would be the time.*

Enter *supercloud*.

[Wikibon has defined it as](#), *“Supercloud describes an architecture, which taps the underlying services & primitives of hyperscale clouds to deliver additional value above and beyond what’s available from public cloud providers. A supercloud delivers capabilities through software, consumed as services; and can run on a single hyperscale cloud or span multiple clouds.”*

Simply put, supercloud is what comes *after* the public cloud: **A value-add layer that hides cloud complexity, exploits differentiated cloud service capabilities, and adds new business (and business network) leverage beyond the raw “Lego™ bricks” of existing public and private clouds.**

Although it represents an evolution beyond just individual public cloud providers, supercloud itself is going through its rapid series of transformations. Looking backward (and forward), we can see it as three eras:

Primordial Era (aka The Age of Terraform) – Common tools enable developers and operators to control disparate and unrelated infrastructure across public and private clouds. Vendors focus on papering over the “control APIs” of public clouds. Value is limited, as data and compute remain wildly disparate and untamed.

Dinosaur Era (aka The Age of Kubernetes) – A common standard for container images enables a “Do-it-yourself” form of supercloud for companies possessing

sufficient expertise and time to build and deploy common containerized capabilities on least common denominator public and private cloud servers. Vendors focus on a “deploy many” approach, where multicloud is achieved manually through redundant construction of similar (container-based) IT stacks on different clouds. Value is higher than for tooling alone but still limited in scope due to the cost and complexity of the approach.

Rise of Humans (aka The Age of SuperServerless) – Platform capabilities delivered in SaaS fashion turn multi-cloud challenges into “just another feature”. The challenges of babysitting infrastructure and its operation are delegated to the cloud hyperscalers, while a new generation of vendors creates a value-added layer for businesses, hiding cloud complexity. **Established vendors such as Snowflake and startups like Vendia offer multicloud and multiparty data and compute as first-class features, enabling applications and business partnerships to span public cloud boundaries without new infrastructure investments.**

Identifying Supercloud Vendors

Supercloud represents the natural democratization of the current “wall of complexity” facing both developers and businesses. They know they need to avoid lock-in, but rolling their multicloud strategy – both at a portfolio level and individually as a feature *enabled for every* one of the hundreds of millions of business applications in operation today – is mind-blowing in its complexity and cost. Public cloud hyperscalers, though they possess the technological and personnel resources to attack the problem, are structurally disinclined to do so: Their balance sheets look better the higher their “garden walls” grow, keeping data, compute, and other clouds spend on their individual platforms. So the natural outlet for this new supply/demand intersection will shift to a new generation of companies, “born in the cloud”, that inherently understand both how to build their infrastructure in a multicloud fashion but also (and more importantly) *how to deliver business value to customers in a cloud-agnostic way.*

Though it’s still early days, those techniques are already becoming apparent:

- **Delivered as SaaS, not infrastructure.** If the route to supercloud requires every developer in every company to be able to design, deploy, and operate every cloud provider, then it’s game over before it starts. Infrastructure in the supercloud era, to hide complexity and offer customers fast time to market, will focus on API-based solutions packaged as SaaS services, not as old-school COTS (Commercial Off-the-Shelf) software.

- **Focused on business networks, not individual accounts.** In an (understandable) effort to limit migration complexity, public cloud hyperscalers focused their initial offerings and service designs on replicating on-prem infrastructure. Each company, each department, each account, and each application becomes an island, creating a series of smaller data and compute silos in addition to the massive public cloud vendor silo itself. The new generation of supercloud offerings focuses on silo *elimination*, with cross-company, cross-cloud, cross-account, and other “gap spanning” features as first-class citizens in every aspect of their designs, from security and authorization capabilities to raw data transfer performance.
- **Designed for solution-oriented value, not low-level DIY.** Public cloud providers have done a fantastic job of creating primitive “building blocks” that make it possible to compose an astonishing variety of outcomes. With the possible exception of high-frequency trading or fully custom supercomputer hardware, anything a company can do itself in an on-prem environment can be accomplished (with enough effort) in a public cloud.

But delivering business value is about time to market and ROI, not just technical feasibility.

Knowing how to assemble the hundreds of services on AWS, let alone make them interact with the thousands of services from other cloud vendors, isn't easy. Supercloud vendors abstract this complexity away. Creating easy, immediate wins for business outcomes from BI and ML (Snowflake) to APM (Datadog) to real-time data sharing (Vendia) enables their customers to focus on their differentiated business problems instead of how to train their developers to compose the ever-changing, low-level public cloud services into useful outcomes.

Managing the coming supercloud migration

As a technical or business leader, how can you make the most of the rise of this new supercloud category of vendors? What best practices can be followed today, when supercloud is still an emerging concept? The following list provides some basic guidance for thinking about cloud investments during the era where both cloud-specific infrastructure and supercloud vendors will both remain critical elements of an IT portfolio:

- **Select supercloud vendors who already “get it”.** The simplest and most highly leveraged way to incorporate change is to select a vendor who can successfully encapsulate it. Unlike hiring an SI for public cloud migration, this strategy aligns well with supercloud’s emergence, because the entire point is that *multicloud is best delivered as a feature*. Moreover, the SaaS delivery and packaging approach of supercloud means that the majority of improvements made by vendors will transparently “flow-through” to their users without the expensive and challenging upgrades of the past.
- **Skip the “uncanny valley”.** Unless you’re an ISV, going it alone – attempting to build complex, cross-company, cross-cloud data and compute sharing into every application – is generally a bad idea. These features are among the most complex in distributed systems design and operation and represent some of the most critical security and compliance exposure among all IT holdings. At the same time, those challenges (and others, like performance and fault tolerance) represent “undifferentiated heavy lifting” – the perfect set of problems to offload to vendors long term.
- **Avoid “accidental multicloud”.** Multicloud strategies and investments are among the most costly and complex in a CIO’s portfolio. Leaving them to chance is perilous, as it can result in costly legacy migrations down the road. Foundational elements, like how analytic and operational data are handled, and burst compute needs, should be viewed through multiple lenses, including how to avoid redundant spending and staffing. Early and intentional supercloud vendor selection can focus those decision processes, reducing complexity and spending over time.
- **Design and invest for network effects, not for individual applications.** One of the most important emerging best practices is to shift focus from the “micro” – each application owned or operated by IT – to the “macro” – the value of automating and digitizing business networks and their associated workflows. Today, [up to 80%](#) of critical business data now lives outside a company’s four walls. This new era of inherently distributed and decentralized applications demands even more from the “IT plumbing” on which it runs, making decisions about multicloud and cross-vendor technology capabilities a business-critical decision *today*. Companies like Vendia turn the mess of multicloud into operational data-sharing solutions, and platforms like Snowflake rein in similar complexity for data lakes and their associated processes.

Conclusion

Supercloud may be yet another tidal wave for technology managers to surf, but it brings with it the promise of taming public cloud complexity and a more mature focus on delivering business value versus just low-level technology innovation. If you would like to learn more visit vendia.net.