CORESTACK®
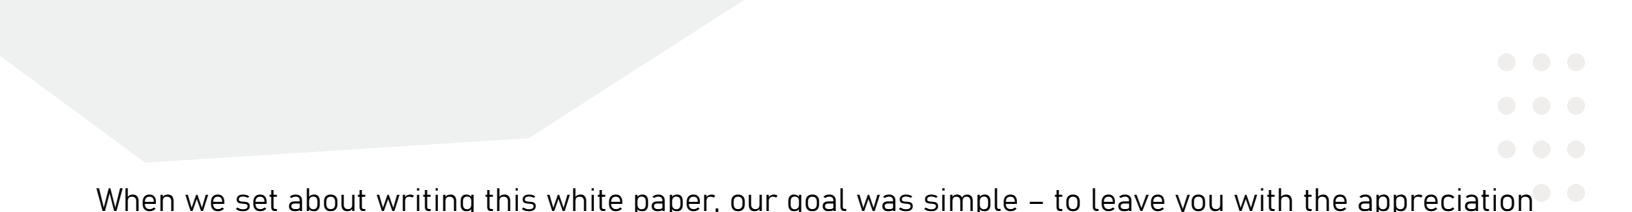
Cloud Governance Meets a Powerful Ally in Generative AI

Authors:
Robert Ford, Chief Strategy Officer, CoreStack
Scott Tucker, Director of Data Science, CoreStack

It's time to unshackle your genius and welcome the smartest person to the metaphorical boardroom – GenAI.

When we set about writing this white paper, our goal was simple – to leave you with the appreciation that, as a leader in the world of cloud and AI, your value will be more in how you think and question than in what you do. After all, the "doing" part of all of our roles will ever-increasingly be in the domain of technology.

Naturally, we will be talking about this topic in the context of cloud governance; after all, CoreStack is an AI-powered next-generation multi-cloud governance platform. However, rather than showcase our AI prowess in an AIOps feature monologue, we want to provoke broader thinking and greater questioning in your mind around something more significant – what's possible when cloud, and more specifically cloud governance, meets generative AI.

There is a lot to unpack in that lofty aspiration, including what we mean by cloud, cloud governance, and generative AI. This white paper will provide CoreStack's point of view on each, bringing these ideas together to ensure that every leader in business is part of the AI revolution, along with – not replaced by – generative AI.

Delivering Value Faster

Let's begin by sharing a cloud point of view. What's most important here is not so much whether you agree with ours, but that you have one of your own. For those asking themselves why, here are three reasons:

- Cloud is about enabling you to deliver value to your customers faster – period. With rapid on-demand resources, infinite scalability and reach, and endless options for experimentation, the cloud can position you to do just that.

- Cloud constructively disrupts everything – which is a good thing. Whether we are talking business models, competitive advantage, architectures, or how we work, there is very little in business that has not been (or will soon be) disrupted by the cloud.

- Cloud is at the center of all digital ambition. Cloud unleashes the art of the possible for every company to exploit if they can front the innovation wave. You just need a visionary purpose that orbits the cloud as its center of gravity – in an orderly manner.

Doing cloud right is critical today, as it is one of tomorrow's three industry-recognized disruptive forces – along with AI and quantum. Let's upgrade this cloud viewpoint into a broader cloud world-view. Cloud explosion continues in all flavors, whether public, private, hybrid, edge, multi-cloud, or even the nascent super-cloud that is gaining attention. According to McKinsey, choice, complexity, and value realization are everyday concerns for the roughly 85% of organizations that have more than one cloud services provider – and for the almost 50% that have three or more.

While the approximate numbers change daily, when you factor in that just one of the big-four cloud hyperscalers offers over 400 cloud services, over 4,000 cloud resources, over 50,000 cloud SKUs, over 5,000 cloud metrics, over 2,000 cloud policies, more than ten cloud frameworks, and support for over 500 cloud standards, it would appear that the potential of the cloud is only limited by your imagination, vision, and confidence.

Fronting this cloud innovation wave demands a world-class approach to optimization, compliance, security, and actual value creation – the kind that speaks to "creating a new dollar" rather than simply saving or stretching an existing one. Here's where we welcome cloud governance to the conversation.

Confidence Matters

Cloud governance is the new frontier of cloud computing that you need to master as you jump headfirst into the world of generative AI. You want to embrace AI with confidence – the confidence of knowing that the cloud powering it is the best it can be based on any vector of governance that matters. This level of confidence requires more than native hyperscaler consoles, homegrown solutions that are poorly stitched together, or siloed third-party cloud management tools.

The bar has been raised from good, independent cloud management to great, interdependent cloud governance. With this level of confidence in the cloud, the boardroom can unleash generative AI into the DNA of their business. Put simply, cloud governance enables you to embrace the cloud with confidence.
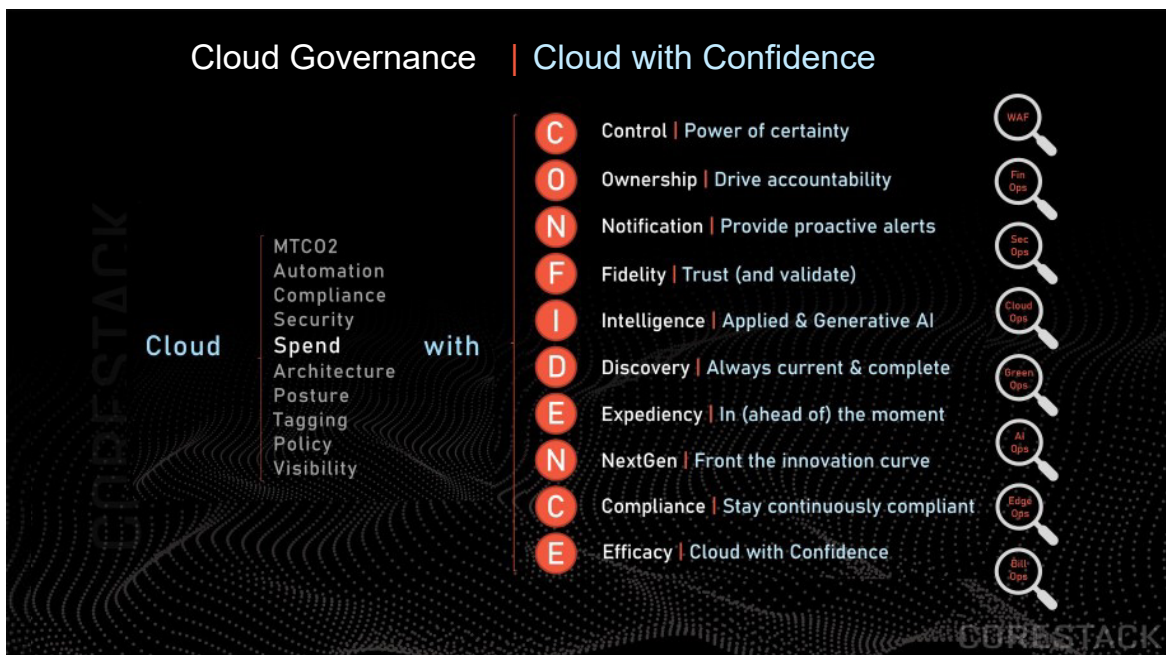
Unprecedented Opportunities with Generative AI

Generative AI is revolutionizing business, from automating mundane tasks to aiding in complex decision-making processes. Its potential economic impact is staggering, with estimates suggesting it could add trillions to the global economy. This transformative power is not just about automation but also about creating new ways of thinking, problem-solving, and even creating. While the printing press democratized information, GenAI has the potential to democratize intelligence and creativity, offering tools that can think alongside us, thereby amplifying human potential.

Expanding the Role of Semantics in Generative AI

GenAI is increasingly being used to infer semantics from a variety of data types, including text, images, diagrams, and graphs. Knowledge graphs, for instance, are being employed to capture semantic relationships, thereby enabling AI models to generate more contextually relevant and explainable outputs. Whereas pretty much every industry these days relies heavily on data inter-pretation, with generative AI it is so much more than "just" data interpretation. It is also the conver-gence of cloud (hyperscalers that make it possible), data, and the variety of ways in which data could be gathered, normalized, processed, analyzed, and interpreted to provide insights. By mastering the art of inferring semantics from diverse data types, generative AI is not only making machines more intelligent but also making them better collaborators in decision-making processes.

Turbocharging Business and Applications

Data comes in many forms – tables, graphs, unstructured text, and more. Transforming this hetero-geneous data into natural language descriptions can help unlock immense business value. Putting information into narrative form allows both humans and AI systems to recognize patterns, relation-ships, and meaningful insights that may not be obvious just looking at the numbers. From a large dataset, key attributes and findings can be extracted and explained in a variety of ways, including writing, images, video, music, and more. This language can then be leveraged by generative AI to automatically produce software designs, applications, and even synthetic code.

Language bridges the gap between raw data and valuable business assets, and helps interpret and generate insights. Those who understand how to leverage the incredible versatility of language will derive the most value.
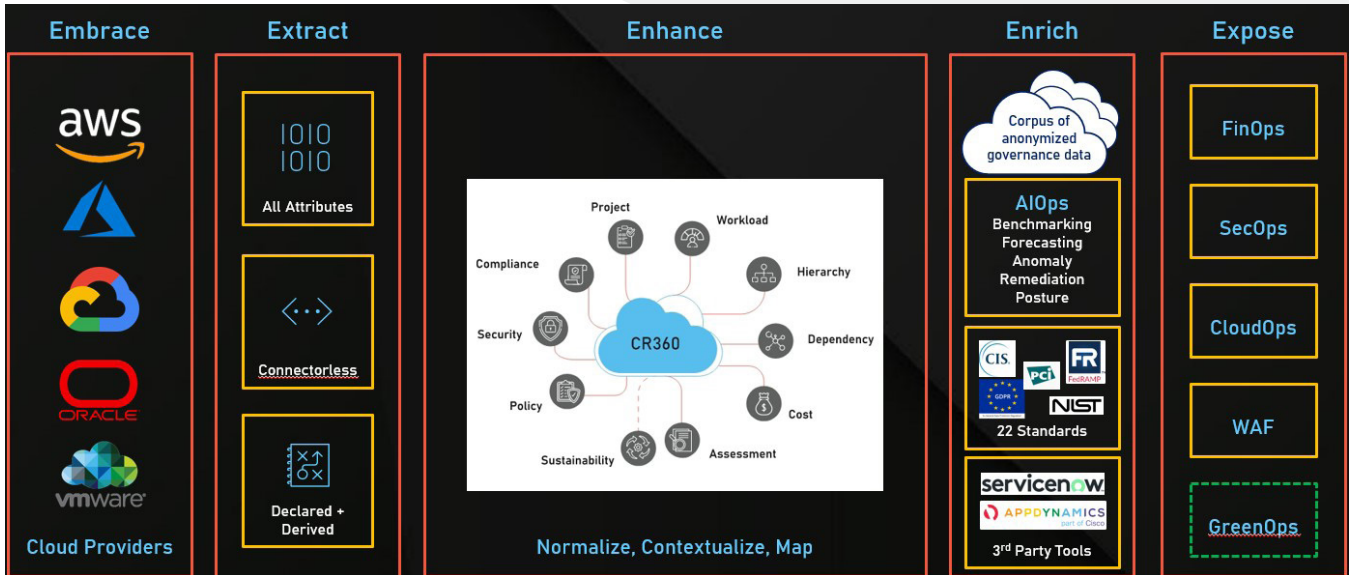
Surviving the Revolution

The rate that generative AI is learning and improving is accelerating. More and more, roles occupied by humans are being taken over by AI. In the past, most of us didn't have to worry about software replacing us, but that has changed; apps can now "think" and possess agency. This shift demands a fundamental change in the way the workforce sees itself:

- AI can be a powerful ally or an unbeatable rival.
- Thinking, imagining, and questioning are emergent new ways of working.
- Collaborating with conversational AI is both essential and a survival skill.
- Using AI to get more out of AI ("AI-on-AI") is an unprecedented force multiplier.
- Semantics has moved from the Ivory Tower to the workplace.

Generative AI Meets Cloud Governance

In order to leverage generative AI in cloud governance, there must first be deep domain knowledge about cloud and multi-cloud computing. There must then be a comprehensive understanding of what "governance" means and what a cloud governance model requires. This knowledge should then connect the hyperscalers, the variety of data from cloud resources, the relationships between those resources, and the interpretation of the data. It should reflect how cloud resources can integrate and connect with relevant backend systems, and how all of these insights can be made available to provide business value.

The better this model of cloud governance is defined and implemented, the more likely you'll be able to harness the unique capabilities of AI and apply them to solve complex business problems. Fortunately, CoreStack already has such a model. It is called CR360™ or Cloud Resource 360, a comprehensive view of multi-cloud resources.

Embrace — **Extract** — **Enhance** — **Enrich** — **Expose**

Embrace
aws
Azure
Google Cloud
ORACLE
vmware
Cloud Providers

Extract
1010 1010
All Attributes
< · · >
Connectorless
Declared + Derived

Enhance
Project
Workload
Compliance
Hierarchy
Security
CR360
Dependency
Policy
Cost
Sustainability
Assessment
Normalize, Contextualize, Map

Enrich
Corpus of anonymized governance data
AIOps
Benchmarking
Forecasting
Anomaly
Remediation
Posture
CIS PCI FR FedRAMP GDPR NIST
22 Standards
servicenow
APPDYNAMICS part of Cisco
3rd Party Tools

Expose
FinOps
SecOps
CloudOps
WAF
GreenOps

The Unique Capabilities of Generative AI

ChatGPT surprised everyone, including its creators, with its unanticipated capabilities. We have spent nearly a year in discovery of its unique capabilities and what it takes to use them reliably and effectively for cloud governance. Avoiding the anthropocentric arguments about the differences between machine–based and "real" (human) intelligence, we can safely say that the unique capabilities of generative AI have the ability to:

- Intelligently communicate with humans via language, diagrams, and other modalities.

- Reason about most things humans can: problem solving, planning, analyzing, synthesizing, interpreting, etc.

- Perform most software engineering tasks: designing, coding, refactoring, documenting, testing, versioning, etc.

- Learn from examples and learn by analogy.

- Create: documents, essays, images, videos, etc.

With CR360 and E^5 (Embrace, Extract, Enhance, Enrich, and Expose) as a foundation, we can continue with the next evolution of cloud governance by introducing machine-based intelligence. This provides two significant advancements:

- **Enhancing a-priori functionality with dynamics.** This means that functions can be created on demand to satisfy customer needs. There is no need to anticipate and enumerate functions and workflows ahead of time.

- **Adding full-duplex interactions with the system.** This means follow-up questions and ad-hoc multi-step activities can be executed conversationally, without having to know which screen or form to pull up next.

Generative AI Applied to E^5

In this section, we highlight the ways generative AI will be applied to CoreStack's E^5.

Embrace

At CoreStack, we use generative AI to create models that capture the essence of the disparate offerings from the various hyperscalers. These models bridge, abstract, and normalize the terminological gap between them, providing the context that allows AI agents to reason about cloud governance free of superficial differences.

We discover all the cloud resources, we collect the billing, compliance, and security information, and we gather the properties and relationships between the cloud resources to create an ontology, mapping, and master data abstraction of all the cloud resources from all the hyperscalers, "embracing" all they offer.

This approach enables customers to engage with AI to compare alternatives, explore optimizations, troubleshoot, and more, conversing with a virtual expert in cloud computing on any cloud governance topic of interest.

Extract

As new services are introduced, there will be new data to capture. CoreStack's patented "connector-less" system and rapid data extraction method for interoperable DSL allow for orchestration of multiple cloud platforms and services. In this environment, generative AI can be used to:

- Dynamically determine the nature of the data.
- Dynamically extract the attributes and metadata as needed.
- Dynamically keep up with the hyperscalers as they evolve and provide newer and more numerous types of services and resources in the cloud.
- Dynamically integrate the new data into CR360.

Enhance

Besides providing dynamic functions for new data, GenAI can provide deeper contextualization. For example, the data can be integrated in the generalized model, making the model more relevant to customers.

Deeper contextualization also means taking advantage of GenAI's ability to use transfer-learning to bring knowledge from a related area to bear. For example, when trying to optimize operations within a budget, GenAI's knowledge from the field of Operations Research and Optimization Theory can be used to provide more options and guidance to the customer, while Petri Nets can model and simulate cloud operations processes, enabling the identification of inefficiencies, bottlenecks, and potential points of failure in cloud workflows.

Another example could be more cloud security related. Here Computational Graph Theory can be used to unravel the complex web of interdependencies within the cloud infrastructure, facilitating better security planning and risk assessment.

Finally, a cloud compliance example, where constraint satisfiability and formal logic can be used to rigorously define and automatically verify cloud compliance by expressing regulatory requirements and cloud configurations as formal constraints, enabling precise detection of non-compliance issues. This approach significantly reduces the risk of violations and enhances the overall security posture.

Enrich

Generative AI's reasoning and language understanding capabilities allow it to rank in the top 10% for many standard measures of intelligence. Combined with its code generation capability, GenAI can answer very complex questions. For example, given a set of log files, workflow descriptions, and operational data from CoreStack, and asked to find patterns in the data, it might generate and execute machine learning code to cluster the data, perform regression analysis, and find correlations within the data.

The real power comes when the results are returned. In conversational mode, the user can ask for explanations in language that is relevant for them, including mini-tutorials, co-pilots, and custom GPTs. The user can ask why certain decisions were made, what the alternative approaches are, what the trade-offs are, and so on.

Here are some more examples of questions that can be used to enrich the data:

- I see that the major vendors have more AI and ML services than any other kind.

    o Why is that?

    o How can I use AI services to improve my use of cloud resources?

    o How can I use AI services to improve my on-premises operations?

- I have uploaded my disaster recovery plan.

    o How can I best use CoreStack's policy and compliance framework to realize my goals?

    o Are there any areas where I should improve in my plan?

- I'm interested in container technology.

    o Would any of my workflows benefit from using containers?

    o I don't know much about them – what are the risks and benefits?

Expose

The primary way to interact with generative AI is through a conversational interface. This interface not only facilitates text communication but also supports the display of various visual elements such as images, diagrams, graphs, and other graphics. GenAI can be used to take exposed data to a whole new level by being able to conversationally display it any way you'd like it.

In February of 2024, OpenAI released Sora, a text-to-video model. According to OpenAI: "We're teaching AI to understand and simulate the physical world in motion, with the goal of training models that help people solve problems that require real-world interaction."

What this means is that very soon you will be able to see simulations of your cloud computing as video unfolding in time. When this capability becomes available, it will be a simple matter, given CR360 and CoreStack's abstract models of cloud computing, to connect your operational data to Sora (or similar text-to-video models) for 3D animations of your operations.

In addition, domain-specific, highly contextualized custom AI agents provide up-to-date information and intelligent assistance to serve every customer need. These agents can integrate with other agents and can evolve as state-of-the-art AI progresses.

Cloud Governance Meets AI-on-AI (CR360 + AI²)

The key to maximizing the value of AI is what we call AI-on-AI or AI². Because GenAI understands language, output from one AI can be fed to another AI to enhance, transform, or interpret business problems and agendas.

CoreStack's AI Agents, or co-pilots, are designed to take advantage of this feature. AI² represents an exponential leap in efficiency and innovation. By enabling AI systems to interact and learn from each other, we unlock new levels of automation, insight, and adaptability that were previously unthinkable. This synergy not only streamlines complex cloud governance tasks but also pioneers novel approaches to managing digital ecosystems and accelerating digital transformation efforts.

Through continuous feedback loops, AI^2 elevates cloud governance, turning data into actionable intelligence at an unprecedented pace. As these intelligent agents evolve, they pave the way for more intuitive, self-optimizing cloud environments, ensuring that governance strategies are not just responsive but also anticipatory, enabling any and every organization to unleash the power of cloud on their terms.
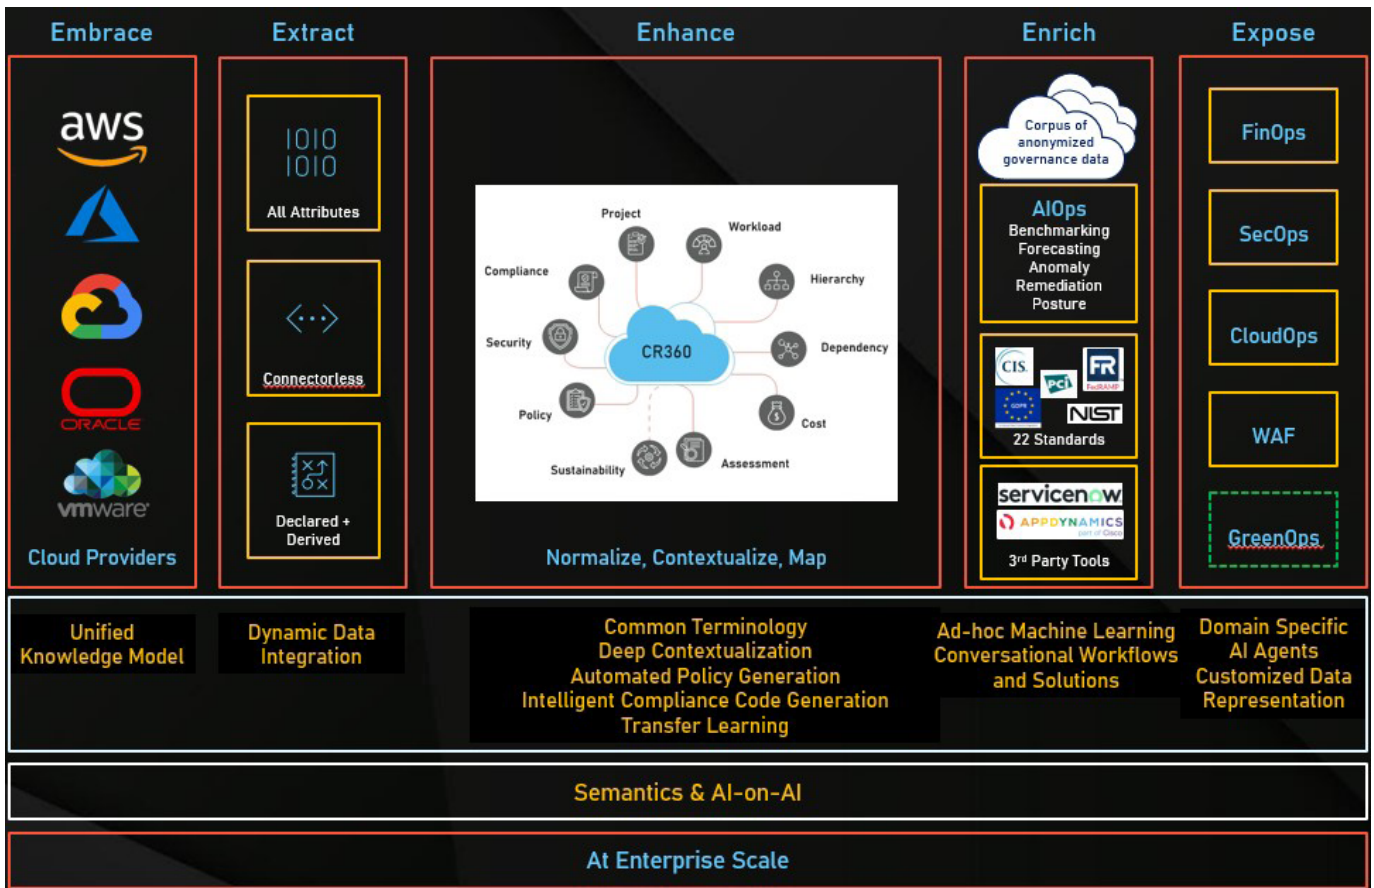
By adding active, deeply contextualized intelligence to every facet of E^5, we bring cloud governance into the age of AI.

A Powerful Ally

AI is not something you buy; it's something you buy *into*, and it can be a powerful ally. Having a solid foundation of cloud governance will enable you to embrace AI and cloud with greater confidence, so you can go further, faster. Hopefully, this white paper has left you with as many questions as answers on the potential of generative AI when combined with world-class multi-cloud governance. We wanted to leave you more informed but also still contemplating what's possible when you leverage a platform like CoreStack that abstracts the complexities of multi-cloud governance and integrate the power of AI in all its forms – be that applied, generative, or whatever comes next.

We've explored CoreStack's point of view on cloud, cloud governance, and generative AI, bringing these exciting, constructively disruptive forces together to hopefully inform and excite you. We want every business leader to be an active participant in the AI revolution, not a casualty of it.

To learn more about CoreStack, visit www.corestack.io or reach out to us at contact@corestack.io.

CORESTACK®

✉ contact@corestack.io

🌐 www.corestack.io

CoreStack is an AI-powered NextGen Cloud Governance platform that enables enterprises to embrace cloud with confidence, rapidly achieving continuous and autonomous cloud governance at scale. CoreStack helps 750+ global enterprises govern more than $2B in annual cloud consumption. The company is a Microsoft Azure (Legacy) Gold Partner, Amazon AWS Technology Partner with Cloud Operations Competency, Oracle Cloud Build Partner, and Google Cloud Build Partner.