Why DoD Should Adopt a Multi-Cloud IT Strategy

By Marcia G. Madsen, Peter O. Schmidt, Luke P. Levasseur, David F. Dowd

Note from the Editor:

This article describes how companies approach cloud based services, and it argues that, in keeping with statutes that mandate competitive procurements, the Department of Defense should follow the lead of the commercial marketplace by adopting a multi-cloud strategy.

The Federalist Society takes no positions on particular legal and public policy matters. Any expressions of opinion are those of the authors. Whenever we publish an article that advocates for a particular position, as here, we offer links to other perspectives on the issue, including ones opposed to the position taken in the article. We also invite responses from our readers. To join the debate, please email us at info@fedsoc.org.

• Jared Serbu, *Pentagon: Need for speed justifies single-award approach to JEDI cloud contract*, Federal News Radio (May 15, 2018), https://federalnewsradio.com/defense-main/2018/05/pentagon-need-for-speed-justifies-single-award-approach-to-jedi-cloud-contract/.

• Amber Corrin, *The Case for One Giant, Multibillion-dollar Cloud Contract for DoD*, C4ISRNET (April 23, 2018), https://www.c4isrnet.com/it-networks/cloud/2018/04/23/the-case-for-dods-single-award-cloud-contract/.

• Alexander Rossino, *Market Analysis: Why Industry Should Not Worry about the DoD's JEDI Cloud* (May 23, 2018), https://iq.govwin.com/neo/marketAnalysis/view/2801?researchTypeId=1.

• Anthony Capaccio, *Pentagon Eases Secrecy Over Cloud Contract as Amazon Rivals Fret*, Bloomberg (May 14, 2018), https://www.bloomberg.com/news/articles/2018-05-14/pentagon-defends-cloud-contract-rivals-call-a-lock-for-amazon.

About the Author:

Marcia G. Madsen, Peter O. Schmidt, Luke P. Levasseur, and David F. Dowd all practice law at Mayer Brown LLP, which represents Oracle Corporation.

Technology changes at the speed of light. Twelve years ago, there was no such thing as an iPhone or android mobile device. When the iPhone was introduced in 2007, Blackberry held a dominant position in the mobile communications market; it has less than 0.1% of the market today. Ten years ago, there was no iPad, Alexa, Uber, Instagram, Snapchat, Kickstarter, or Square. And although the idea of networked computing traces its lineage back to the 1960s, the term "cloud computing" wasn't coined until 2006—and it is just within the last ten years that commercial cloud-based services and storage offerings have exploded. In light of this pace of development, no reasonable consumer—or large IT buyer—would lock itself into a single technology or service as its exclusive choice for the next decade.

In early March 2018, the Department of Defense (DoD) issued a draft request for proposal (RFP) for the Joint Enterprise Defense Infrastructure (JEDI) contract. It issued an amended JEDI RFP on April 16, 2018, and a final solicitation may be issued later this summer. The JEDI procurement involves cloud computing infrastructure as part of DoD's effort to modernize its IT services. JEDI will apply across DoD and is valued in the billions of dollars. JEDI is not an acquisition for a single, broad-ranging cloud project, but is for an indefinite delivery, indefinite quantity (IDIQ) contract vehicle, under which work will be parceled out through separate orders for different requirements (and at different security levels) as particular agency needs are identified.[1]

The JEDI procurement has received substantial attention—and been the subject of vigorous debate—because of size of the contract (estimated at $10 billion) and the projected size of the government cloud services market. Virtually all knowledgeable analysts recognize that DoD will experience significant efficiency (and other) gains from migrating much of its IT services to cloud-based facilities.[2] One of the most intensely debated aspects of JEDI is whether DoD should migrate applications and storage to a single cloud service provider (CSP), or multiple CSPs.

DoD has asserted that it intends to award JEDI to a single vendor. As we explain below, relying on a single vendor for JEDI cloud services would be a serious and unnecessary error. As the commercial marketplace demonstrates, multiple cloud solutions reduce enterprise costs, increase agility, insulate customers from problems from a single point failure, and offer substantial performance and security benefits. No evidence-based or coherent explanation for selecting a single-cloud provider for JEDI has

1 *See* FAR 16.504(a).

2 *See* Phil Goldstein, *DOD, State Department See Benefits from Shifting Global Operations to the Cloud*, FedTech (July 14, 2017), https://fedtechmagazine.com/article/2017/07/dod-state-department-see-benefits-shifting-global-operations-cloud; *see* Joint Enterprise Defense Infrastructure (JEDI) Cloud DRAFT Statement of Objectives (SOO) (April 16, 2018), https://www.fbo.gov/utils/view?id=09e2b02eb15b49a0b4e37ad121dbee3c, at 1 (draft JEDI RFP, Statement of Objectives, explaining that, among other things, large-scale migration to a cloud is necessary to avoid "environments [that] are not optimized to support large, cross domain analysis using advanced capabilities such as machine learning and artificial intelligence to meet current, and future 17 warfighting needs and requirements").

been provided, and doing so will stifle innovation and increase government costs.

Limiting the JEDI procurement to a single CSP's technology solution also cannot be reconciled with a statutory requirement for multiple awards for contracts that will involve the issuance of orders for a period of time after award.[3] That requirement was implemented to give the government the benefits of innovation and price that can be achieved through periodic competitions among a qualified group of suppliers. For the substantial JEDI cloud acquisition, however, DoD appears intent on ignoring the innovations available in terms of technical merit and price—and the other expected benefits of an ongoing competitive environment for which Congress has expressed its statutory preference.

This procurement has aroused much interest because industry understands that cloud computing is likely to develop into a massive and rapidly changing market. As with other technology fields, new entrants are appearing frequently, and leading-edge capabilities are changing rapidly. As a result, in mid-2018, it is not possible to predict which technology or CSP or approach will be the most capable in three years, or which company might provide the best pricing in two years. Given this changing environment, the government should put itself in a position to take maximum advantage of innovations in the market and not tie its hands. But DoD is ignoring the manner in which commercial buyers are reacting to the changing cloud services environment (and the recommendations of industry) and, instead, apparently intends to proceed with a single-vendor approach that will preclude any further consideration of options for a particular application (under JEDI).

It is possible that DoD will do well by locking itself into a single company's technology. But given the pace of technological change, the time and money at stake, and the statutory preference for the flexibility of a competitive multi-vendor IDIQ contract vehicle, that is a lot of eggs to place in one basket.

I. Background of Cloud-Based Services

A. The Industry Push for Ecosystems

Certain companies within the IT industry have touted mass adoption of a single company's technology ecosystem as the most efficient way to implement IT solutions.[4] Although a common platform can increase certain efficiencies—such as having a single point-of-contract for customer support—walled-off ecosystems result in high exit barriers and make expanding into other product lines of the same company the path of least resistance. Such concerns have arisen with respect to Cisco's Enhanced Interior Gateway Routing Protocol and Apple's products in its consumer market. The common platform sales pitch has now metastasized into the cloud computing market.

A company that can garner sufficient critical mass to exercise some control on a market can experience substantial rewards by promoting an ecosystem solution. To counter this strategy, competitors have evolved an approach that involves pushing cross-platform integrations. Such integrations allow individual participants to develop a best-of-breed technology in one area and integrate with the products of other companies (which may be best-of-breed in other areas). The resulting hybrid solutions can outperform any single vendor's ecosystem.

Although various participants in the cloud computing market have pushed the single ecosystem concept, the structure of the market is contrary to a single-cloud environment. In part, that results from the fact that, unlike earlier IT markets, the best cloud-based services have been built from the ground up with the intention of integration. That intention is fundamentally thwarted by a walled-off system that makes integration across platforms difficult or impossible.

B. What Is Cloud Computing?

Generally speaking, cloud computing is a business model for renting access to shared software and hardware over a remote network—usually the internet. Although there are many nuances to cloud offerings, they are generally separated into Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS).[5] Most users are familiar with SaaS functions such as creating a Google document in a web browser. In this case, the software managing the document is not running locally on the user's computer, but remotely on servers owned and operated by Google. Such cloud-based services can be most sharply contrasted with on-premises solutions in which the user controls both the IT hardware—including acquiring, locating, maintaining, servicing, and repairing that hardware—and the software running on the hardware.

With IaaS, a cloud provider offers virtual hardware. The customer interacts with the hardware over the internet just as it would with a physical (albeit remote) server. For example, the customer might install a new web server on its IaaS, and then use that software (housed on cloud-based hardware) in the same manner as software housed on hardware at the user's facility. PaaS refers to services that offer software resources, such as database software or a development environment, that developers can use to build more customized applications.

In providing services to customers, a CSP scales its servers up or down to meet demand. A customer will interact with the CSP through a single interface, while the CSP may use multiple physical servers (typically shared with other customers) to deliver the required computing and storage power. With basic computing resources such as storage, processing power, and network bandwidth available in ever greater supply through cloud computing services, prices have been falling, increasing the number of customers able to access and afford large scale computing systems.

3 10 U.S.C. § 2304a(d).

4 A cloud ecosystem is a complex system of interdependent components engineered to work together to enable cloud services. *See* Definition: Cloud Services, TechTarget, https://searchitchannel.techtarget.com/definition/cloud-ecosystem.

5 *See* Connor Forrest, *SaaS, PaaS, and IaaS: Understand the differences*, ZDNet (Nov. 1, 2017), https://www.zdnet.com/article/saas-paas-and-iaas-understand-the-differences/.

C. Business Benefits of Different Cloud Environments

Commercial businesses and governments can choose to migrate applications and storage to only a single cloud or to use multiple cloud service providers (multi-cloud). Multi-cloud offerings provide at least four important benefits.

First, multi-cloud arrangements limit vendor lock-in and thereby increase cost competition. It is relatively easy to build a solution that is independent of any specific provider (and can therefore operate in multiple clouds) if this need is considered during the design phase. Unfortunately, naïve customers often find themselves with provider-specific solutions and, as a result, have little leverage in future negotiations (because they are locked in). Although a customer may have received bulk discounts and other cost savings up front, the vendor lock results in the customer being unable to take advantage of price reductions that result from the constantly declining cost of technology and the relentless commodification of cloud services. By pursuing a multi-cloud strategy from the outset, customers set themselves up for long term savings.

Second, the use of multiple CSPs allows customers to be agile and ramp up or down usage of a provider based on factors such as features, performance, and cost. For example, no one CSP will be the best on every metric in every region. Low latency may be important for a particular application in a particular region. But waiting for a certain performance level to be achieved by a CSP in a given region can compromise effectiveness. The ability to use multiple providers greatly increases coverage of resources (wherever located). Accordingly, customers are able not only to achieve target performance as early as possible, but they can effectively leverage price competition when other providers later build equivalent capability.

Third, a multi-cloud approach insulates customers from catastrophic failures by a single provider. With one cloud, a systemic failure by the provider (such as an unpatched security vulnerability or a design flaw) can result in substantial parts of the customer's cloud system becoming inoperable. For example, during Amazon Web Services' (AWS) February 2017 S3 storage outage, connected systems failed as well. A company's or government's use of multiple CSPs introduces redundancy that limits any failures to an isolated component or subsystem. Every provider's systems are subject to human error and failure. Multi-cloud environments reduce the likelihood of system-wide problems as a result of such errors and failures.

Fourth, using a range of providers enables customers to take the path of least resistance for migrating each application. For example, it is generally faster and cheaper to migrate an Oracle on-premises database to an Oracle cloud database, as opposed to another CSP's cloud database. Assuming several providers would perform equivalently for a particular application, significant savings can be achieved by this approach.

Very few large businesses opt to use or pursue a single-cloud environment for their applications. Customers in the CSP market have found that the principal theoretical benefit of having all applications and storage on a single cloud—i.e., fewer consistency problems resulting from the data being consolidated in one place with a single architecture—are less than the substantial benefits of having best-of-breed technologies from different offerors.

D. Technical Benefits of Multi-Cloud Environments

1. Performance

One of the principal benefits of multi-cloud is the ability to use best-of-breed implementations regardless of which entity developed a given implementation. This is facilitated by microservice architectures, in which applications are structured as collections of independently deployable services that communicate with each other. The same interfaces that allow developers to integrate a provider's services into their workflows can be used to interact with other CSPs' services. Examples of this capability can be found in "serverless" compute, such as AWS Lambda or Azure Functions, in which workloads are triggered by "hypertext transfer protocol" (http) webhooks or callback functions (whether from another provider, IaaS compute, or a customer's on-premises service). Microservice architectures, plus the extensive integrations available, allow best-of-breed selection on a service-by-service basis. As a result, well-designed systems can use the best service available for each function.

There are no substantial performance drawbacks to this approach that would favor the alternative use of a single-cloud environment. Indeed, microservice architectures are best practice even with a single CSP because they facilitate continuous delivery, in which software can be released to production at any time, enabling changes to be rolled out on much shorter timescales (e.g., hours or days instead of weeks or months). In addition, the speed of inter-cloud communication can be very fast because cloud service providers have high speed connections to other providers.

2. Security

In addition to encryption and identification protocols, an important way companies (and CSP customers) improve security in cloud environments is by segmenting workloads across services and providers. This works based on the shared security model of working with CSPs. In short, CSPs certify that they provide security at the standard applicable to the systems for which they are responsible (e.g., from the host server's hardware up to operating system). The customer is then responsible for ensuring that its data is secure the rest of the way. Thus, as with on-premises solutions, the customer must maintain security domains and separation of access. This provides the opportunity to segment workloads across services and providers, which further secures operations.

Even businesses or governments with a single cloud provider generally segment workloads, e.g., across multiple accounts and locations. Such segmentation not only allows for complex workloads to scale and to be broken down into manageable parts, it also limits the impact of system failures, human mistakes, and security events. The use of multiple providers further limits such potential problems because the systems are further separated. For instance, different providers have different security models. Enforcing logical access control (i.e., the tools and protocols used for identification, authentication, authorization, and accountability) at the provider level therefore makes it harder to grant accidental or unintended access to services that exist on a separate provider.

3. Support for Different Environments

Although using multiple providers for similar functions creates some concern regarding added complexity, an entire industry has developed within the cloud marketplace in response to the need to support multi-cloud environments. Tools to address complexity range from very specific to very broad in scope and cover everything from cost tracking and analysis to infrastructure and software management. It is easy to find support for any of the major CSPs. And because these tools are often focused on a specific issue, in many cases they address that issue better than the native tools of one provider. This is apparent in many budgeting and cost control tools, in which the graphs and reports are extensive and extensible.[6] Likewise, the tools, such as Terraform for managing infrastructure, often are easier to use and provide a more complete picture of a customer's overall system as compared with native tools.[7]

These kinds of tools are not generally required to help single-cloud environments, which can result in upfront savings and faster procurement. However, these savings can be offset by ecosystem lock-in and worse performance from the inability to use best-of-breed solutions. In this regard, single-cloud environments may be better for small, quick-and-dirty projects, but are generally inferior for large, complex, long-lifetime projects.

II. The Legal Standard Applicable to IDIQ Procurements

As explained above, DoD has issued the JEDI procurement as a single-award IDIQ. However, procurement law favors competition.[8] That preference for competition extends to the award of IDIQ contracts. Dating back to the mid-1990s, there has been a statutory preference for the award of multiple IDIQ contracts rather than a single award.[9]

The preference for multiple IDIQ contracts has been strengthened over time, as the law now prohibits a single award unless the agency makes a mandatory formal determination that a single-award IDIQ (in which there will be no competition for orders) is advantageous for the government. Specifically, statutes applicable to DoD make clear that IDIQ contracts valued over $112 million may not be awarded to a single source unless the head of the agency determines in writing that:

> (A) the task or delivery orders expected under the contract are so integrally related that only a single source can reasonably perform the work;
>
> (B) the contract provides only for firm, fixed price task orders or delivery orders for—
>
> (i) products for which unit prices are established in the contract; or
>
> (ii) services for which prices are established in the contract for the specific tasks to be performed;
>
> (C) only one source is qualified and capable of performing the work at a reasonable price to the government; or
>
> (D) because of exceptional circumstances, it is necessary in the public interest to award the contract to a single source.[10]

JEDI encompasses a variety of cloud services to be implemented through various tasks across DoD at various classification levels, including unclassified work. Given the prevalence of the commercial cloud computing market, it is simply not credible to assert that only one company could perform such work at a fair and reasonable price.

Federal Acquisition Regulation (FAR) Part 16 implements the preference for multiple awards. The regulation provides that the contracting officer must, "to the maximum extent practicable, give preference to multiple awards" of IDIQ contracts.[11] It emphasizes that each awardee "need not be capable of performing every requirement as well as any other awardee under the contracts."[12] Thus, the FAR expressly contemplates that contract awardees will be variously situated in terms of capabilities.

The purpose of the multiple award preference is to enable the government to obtain the benefit of recurring competitions for work that cannot be specifically defined initially but can be identified sufficiently with respect to discrete orders. Multiple awards give the government significant leverage. Companies must compete first to ensure they can be among the awardee group. The government then conducts competitions among the awardee group for task orders that likely will involve variation in needs. Multiple awards provide contractors with diverse strengths. Moreover, the contractors are incentivized to provide excellent performance at ever-more competitive prices throughout the contract term. Even the contractor with the best proposal at the contract stage may not have the capability to propose on every potential order, and each contractor must stay on its toes and sharpen its pencil in the successive rounds of competition at the order level. In multiple-award IDIQs, agencies that manage the contract well should hold a strong hand of cards that improves over time.

6 *See* Seamus Holland, *How Cost Analysis Tools Can Prevent Cloud Computing Calamity*, Programmable Web (Oct. 16, 2017), https://www.programmableweb.com/news/how-cost-analysis-tools-can-prevent-cloud-computing-calamity/elsewhere-web/2017/10/16.

7 *See* Piotr Gospodarek, *CloudFormation vs Terraform*, Medium (Oct. 11, 2017), https://medium.com/@piotrgospodarek/cloudformation-vs-terraform-990318d6a7de.

8 *See*, *e.g.*, 10 U.S.C. § 2304, *et seq.*

9 *See* Pub. L. 103-355, § 1004.

10 10 U.S.C. § 2304a(d)(3).

11 FAR 16.504(c)(1)(i).

12 FAR 16.504(c)(1)(ii)(A).

III. DoD's Stated Rationales for Single-Cloud Are Not Supported

On May 14, DoD released a congressional report that advances a number of arguments in favor of a single CSP.[13] DoD's explanation for pursuing a single-cloud IDIQ does not address the legal requirements for such a vehicle, though it indicates that it will release an explanation addressing these requirements at a later date. The decision to use a single award was apparently made without legal analysis, which apparently is being developed after the fact. Instead, DoD attempts to provide a technical explanation for its decision. In doing so, DoD purports to take a cautious approach by proceeding with a single cloud. But its arguments regarding why a single-cloud solution is purportedly cautious are rooted in at least seven crucial misunderstandings of multi-cloud solutions.

First, DoD states that "[r]equiring multiple vendors to provide cloud capabilities to the global tactical edge would require investment from each vendor to scale up their capabilities, adding expense without commensurate increase in capabilities."[14] In other words, DoD is concerned about the possibility of paying numerous vendors to scale-up such that they can provide capabilities worldwide (to the tactical edge), only to obtain duplicative capabilities.

In the commercial marketplace, cloud buyers have found that the use of multiple vendors enables providers to scale up faster and at lower cost. A multi-cloud approach would allow DoD to use different vendors in different use cases or environments where they *already* are optimal—instead of waiting (and paying) for one vendor to scale up in every area necessary. Although it is probably correct that "no other industry sector matches the scale and diversity of DoD's tactical edge needs," DoD recognizes that "certain industry sectors like oil and gas and university research have motivated vendors to develop commercial capabilities that can, at least to some degree, provide cloud computing and storage resources in austere and connectivity deprived environments."[15] Those capabilities already exist across multiple vendors, and adopting a single-provider approach forces that provider to duplicate capabilities it does not already have.

In addition, DoD appears to recognize the market imbalance that would be created by paying only one provider to develop capabilities everywhere (to the "global tactical edge") and thus plans to include in "the JEDI Cloud contract . . . a requirement for the contractor to provide a detailed portability plan."[16] Such a requirement will not ameliorate the imbalance DoD's plan would create. Realistically, after one vendor has built to the global tactical edge (and been paid to do so), it will be difficult, if not impossible, for other vendors to compete.

Second, DoD asserts that "[w]hile security of data within clouds is largely standard and automatic, managing security and data accessibility between clouds currently requires manual configuration and therefore introduces potential security vulnerabilities, reduces accessibility, and adds cost."[17] In fact, an important premise underlying this assertion is incorrect. The security of data within clouds is not automatic: even with a single CSP, the customer must configure the environment and use encryption to fully control its data.[18] Improperly configured environments can expose data, as evidenced by the numerous cases with AWS S3 buckets in the news.[19]

Other security-related concerns, such as access control, do not indicate that implementation of a single-cloud approach is any safer than a multi-cloud environment. Tools such as Cisco CloudCenter manage access control and encryption across cloud environments.[20] Using such tools, there is little difference between managing security for a single cloud or multiple clouds. Moreover, as explained above, segmenting workloads across services and providers *increases* security. Assuming DoD "make[s] extensive use of containerization," "data standards," and "application programming interfaces which expose the data over secure, modern protocols," as it asserts it plans to do,[21] it should be well-positioned to take advantage of this opportunity.

Third, DoD states that "[m]aintaining inconsistent and nonstandardized infrastructures and platform environments across classification levels complicates development and distribution of software applications, potentially adding delays and costs."[22] In other words, DoD believes it will be easier to develop and distribute applications on a single cloud infrastructure/environment, as a single system will be easier for personnel to learn and use—and easier to secure across classification levels.

The commercial market's experience, in which customers manage data across different infrastructures and platforms, cannot be reconciled with DoD's statement that unnecessary delays and costs will be added by using multiple clouds. This rationale is also inconsistent with DoD's portability requirements. DoD acknowledges it "must strive to make applications portable,"[23] which means platform-independent security. That required portability is inconsistent with the notion that there are substantial benefits to requiring use of a single cloud.

With respect to security, CSPs certify that they provide security for the systems for which they are responsible. DoD is

13 U.S. DoD, *Combined Congressional Report* (2018).

14 *Id.* at 4.

15 *Id.* at 9.

16 *Id.* at 12.

17 *Id.* at 4.

18 Under a shared security model (*see* Section D.2), the customer *must* use encryption to fully control its data, whether in transit or at rest.

19 *See, e.g.*, Dan O'Sullivan, *Dark Cloud: Inside the Pentagon's Leaked Internet Surveillance Archive*, UpGuard (Apr. 30, 2018), https://www.upguard.com/breaches/cloud-leak-centcom.

20 *See, e.g.*, Cisco, *Cisco CloudCenter Solution* (2017), https://www.cisco.com/c/dam/en/us/products/collateral/cloud-systems-management/cloudcenter/at-a-glance-c45-737051.pdf.

21 *Combined Congressional Report* 10, 12.

22 *Id.* at 4.

23 *Id.* at 10.

responsible for security the rest of the way. Assuming DoD adopts a set of standards for CSPs and a set of platform-independent standards for application and data security, managing a multi-cloud environment across classification levels should require a small amount of additional work that is outweighed by the benefits of a multi-cloud environment.

Fourth, DoD states that the "[u]se of multiple clouds would inhibit pooling data in a single cloud (i.e., a 'data lake'), limiting the effectiveness of machine learning" and artificial intelligence (AI).[24] This appears to be DoD's principal objection, and it is repeated throughout the report.[25] DoD asserts that "[m]arket research also indicated that initial migration to a single cloud is consistent with industry best practice"; the only rationale it provides to support that assertion is its concern regarding data lakes.[26]

Data lakes are large repositories of raw data in native formats that, with the appropriate storage and processing tools, can be queried by a user. Even if the development and use of data lakes were "best practice"—and they are not—a single cloud is not necessary for a data lake, as companies have stepped in to fill the need for multi-cloud implementations. For example, Cloudera offers software for running multi-cloud Hadoop data lakes.[27]

Contrary to DoD's assertion, data lakes are *not* considered a "best practice" within industry. Industry recognizes that although use of a data lake benefits IT in the short term (as IT no longer has to devote resources to understanding how information is used when it is dumped into the data lake), getting value out of the data remains the responsibility of the business end user—and without at least some semblance of information governance, the lake ends up being a collection of disconnected data pools or information silos all in one place (a "data swamp").[28] Organizations have found that data lakes are expensive and time-consuming to coordinate, build, and maintain.[29] With AI services having rich data source integrations, the cost and challenges of data lakes can be addressed by avoiding data lakes and allowing the AI services to pull from the original data repositories.

Data lakes also pose substantial risk to security and access control that have not been solved and are not addressed by DoD. Many data lakes are being used for data whose privacy and regulatory requirements are likely to represent risk exposure, but without enforced procedures for placement of data in a lake (where security capabilities and technologies have not been fully developed), it is not clear how the security requirements can be satisfied (particularly if left to non-IT personnel). Moreover, DoD does not explain how it will address the risks associated with a single location of data—or how a single point of failure would better facilitate resiliency, as compared to a scenario with multiple such points.

Finally, maintaining a large amount of data in a single CSP's managed database service poses a high risk of lock-in. As one industry observer noted: "[M]any cloud vendors make it very difficult to extract data, configuration artifacts, and key application settings. This means that if rates rise, your freedom of movement is restricted. Even though your data is technically yours, it's under the control and influence of someone else."[30] DoD does not explain how its vision of a data lake with a single CSP avoids the risk of potential lock-in.

Fifth, DoD notes that its experience to date shows that "hundreds of cloud initiatives have created numerous seams, incongruent baselines and additional layers of complexity for managing data and services at an enterprise level."[31] The report then asserts that "[s]cattering DoD's data across a multitude of clouds further inhibits the ability to access and analyze critical data."[32] DoD's concern is understandable. Interoperability is important for many reasons, including efficiency and maximizing data value. But while single-cloud solutions facilitate interoperability in some ways, by far the largest determinants of success with cloud migration are a business' or government agency's internal migration and development strategies. Thus, DoD's experience to date does not represent a failure of multi-cloud; rather, it reflects the lack of a market research, planning, and a consistent strategy for cloud migration and management.[33]

Sixth, DoD argues that it will capture some of the benefits of multi-cloud by contract, stating that "the JEDI Cloud contract will require ongoing commercial parity of technical offerings" and that "contract clauses [will] ensure DoD continues to get the best pricing as global marketplace pressures drive prices down."[34] Certainly, such clauses are better than nothing, and will

24 *Id.* at 4.

25 *E.g.*, *id.* at 5 ("Leveraging ML/AI at a tempo required to be relevant to warfighters, however, requires significant computing and data storage in a common environment."); *id.* at 6 ("The lack of a common environment for computing and data storage also will limit the effectiveness of ML/AI for warfighters."); *id.* at 9–10 ("In addition to having a consolidated data lake, market research makes clear that a well-articulated data strategy, including an architecture and data storage standards, is critical to realizing the benefits particularly with regards to ML and AI.").

26 *Id.* at 9. The report also quotes two Gartner reports, but those quotations only say that the transformation to the cloud should be staged in some way.

27 *See* Cloudera, https://www.cloudera.com/products/cloud.html; *Hadoop*, Datafloq, https://datafloq.com/hadoop/?utm=internal ("Hadoop is a Free Java programming structure" that supports "disseminated applications running on vast groups of thing machines that process enormous measures of data.").

28 Gartner, *Gartner Says Beware of the Data Lake Fallacy* (July 28, 2014), https://www.gartner.com/newsroom/id/2809117; *see* Dan Woods, *Why Data Lakes Are Evil*, Forbes (Aug. 26, 2016), https://www.forbes.com/sites/danwoods/2016/08/26/why-data-lakes-are-evil/#2f0b2baa4f73.

29 James Ovenden, *Say Goodbye to Your Data Lake in 2017*, Innovation Enterprise (Jan. 10, 2018), https://channels.theinnovationenterprise.com/articles/say-goodbye-to-your-data-lake-in-2017.

30 Dan Woods, *Five Ways To Avoid Cloud Lock-In*, Forbes (June 20, 2017), https://www.forbes.com/sites/danwoods/2017/06/20/five-ways-to-avoid-cloud-lock-in/3/#65a553bb5114; *see* Glenn Solomon, *Why Multi-Cloud is the Next Big Thing in Technology*, Going Long (Nov. 6, 2017), https://goinglongblog.com/multi-cloud-next-big-thing-technology/.

31 *Combined Congressional Report* 7-8.

32 *Id.*

33 *See id.* at 7 ("The DoD's adoption of cloud services to date has been mainly decentralized . . . .").

34 *Id.* at 11.

likely provide some measure of savings. However, they are not as effective as a multi-cloud approach. Intellectual property laws and vigorous competition ensure that there is never true commercial parity of technical offerings—no software provider has ever been best-in-class in every area. And although most favored nation clauses (MFNs) can help offset the lack of bargaining power with a single CSP, DoD will be subject to the nuances of the terms, and the vendor will likely have an information advantage. For example, if an MFN clause covers a specific region but no other vendor operates in that region, the clause will have no effect and provide no benefit to the government.

Seventh, DoD acknowledges that "[i]f the commercial cloud marketplace offerings evolve to become interoperable and seamlessly integrated, DoD could have the ability to meet warfighting and business requirements by employing a range of future contract and award types."[35] This is not so much an independent rationale for a single-cloud approach as a restatement that one of DoD's primary concerns is interoperability. But the evolution DoD is waiting for *already has occurred* and, although it is true that single-cloud solutions provide a modest benefit in facilitating interoperability, the best cloud technology has been built from the ground up with integration in mind. With a well thought out strategy and approach for cloud migration, management, and development, and with the use of currently available tools, DoD's interoperability concerns can be fully addressed.

IV. Concerns Related to DoD's Additional Reasons for Its Single-Award Strategy

DoD's congressional report also attempts to defend offering JEDI as a single-award IDIQ by noting that, as currently structured, JEDI "only" calls for the award of a two-year base period, with the remainder of the ten-year term structured as options. Therefore, the government is not locked into a ten-year contract.[36] That is true so far as it goes, but options typically are exercised where there is satisfactory performance by awardees. One thus reasonably should expect the options will be exercised for JEDI.

DoD also argues that JEDI will not be the only source for cloud services, as DoD already has stated in addressing the single-award approach.[37] Although DoD has other vehicles under which it can procure cloud services, it surely intends to place heavy reliance on the JEDI vehicle. There is simply no reason to compete and award such a large contract if DoD does not intend to use it as a resource. Moreover, the argument that DoD will rely on multiple resources (including but not limited to JEDI) is inconsistent with the contention that DoD will be better off with a single JEDI provider.[38]

Finally, the currently planned JEDI has been defended by people arguing that the contract will allow for new services to be added. Although that is correct, multiple-award IDIQs could similarly allow for the addition of other services, and competition would give incentive to add such services at the best possible price. It is not clear why the government has opted to forego the leverage it enjoys with multiple awards.

V. Conclusion

Today, most businesses already take a multi-cloud approach. This enables them to keep costs down, be agile, and insulate themselves from single point failure. It also provides them with enhanced performance (e.g., by combining best-of-breed services) and enhanced security (e.g., by segmenting workloads across multiple providers and enforcing logical access at the provider level).

As explained above, the relevant procurement statutes and the FAR make clear that Congress' strong desire for competition in federal contracting extends to a preference for multiple-award IDIQ contracts whenever possible. Upon examination, the justifications for ignoring that preference and awarding a single IDIQ JEDI contract are not well supported. Best industry practices and the law counsel that DoD should consider this issue more carefully, follow congressional intent, and adopt a multiple-award approach for a large cloud procurement.

35 *Id.*

36 *Combined Congressional Report* 11.

37 *Id.*

38 DoD also argues that it can protect the government's interests by implementing "contract clauses that ensure that DoD continues to get the best pricing as global marketplace pressures drive prices down." *Id.* But why would DoD rely on a remedy that requires pricing disclosures, continuous marketing, and potential after-the-fact remedies in the event the JEDI contractor does not abide by pricing obligations instead of holding simple competitions—and relying on the market to ensure favorable pricing.