Community Cloud Decision Support

Monica M. Waters, Christine O. Salamacha, and Peter P. Pandolfini

ABSTRACT

Government agencies are migrating information technology applications to community cloud environments to achieve efficiencies and increase productivity. However, the rate of adoption has been slower than anticipated, largely due to concerns about security and overall mission assurance. In this article, the authors describe discovery research conducted in 2012–2013 to identify a repeatable approach for making an early suitability assessment of community cloud environments for DoD missions and to possibly help identify policy changes needed to mitigate the potential risks of planned migrations. Although the research team focused on DoD norms, the researchers hypothesized that this proposed method might also be applicable to other sectors looking for quick, early insights when planning migrations to a community cloud environment. DoD's cloud strategy, policies, and standards have evolved and matured in a number of ways since this research was conducted, but the proposed decision method is tailorable to current applications.

INTRODUCTION

Driven by President Obama's Cloud First policy and evolving budget priorities, the Office of Management and Budget and the DoD Chief Information Office (CIO) instituted strategies to accelerate the pace at which federal entities migrate to cloud computing.[1,2] As a result, DoD organizations face the daunting task of moving mission capabilities from internally hosted and managed environments to off-site environments shared with other entities and managed by a DoD or a commercial cloud service provider (CSP). [In this article, we refer to an organization considering cloud computing as a consumer, whether or not that entity is currently using cloud services. The *NIST Cloud Computing Reference Architecture*[3] defines five major actors: cloud consumer, cloud provider, cloud carrier, cloud auditor, and cloud broker. Each actor is an entity (a person or an organization) that participates in a transaction or process or performs tasks in cloud computing. The cloud consumer is the principal stakeholder that uses the cloud computing services. A cloud consumer represents a person or organization that maintains a business relationship with, and uses the service from, a cloud provider.]

The *Cloud Market Maturity Study* (2012)[4] examined the most significant challenges government agencies identified in migrating to cloud computing. The study concluded that an imbalance in perceived risk over promised benefits was influencing (and constraining) cloud adoption. Frequent news stories about sensitive information being exposed from "the cloud" only added to the

anxiety that government information technology (IT) managers had about migrating to a cloud environment.

Against this backdrop, in 2013 the Johns Hopkins University Applied Physics Laboratory (APL) conducted research to develop a repeatable approach for assessing suitability of community cloud environments for DoD missions. This approach, called the Community Cloud Decision Method, is referred to throughout the rest of this article as simply the Decision Method. The Decision Method provides a quick tool for identifying risk and determining policy changes needed to mitigate risk for community clouds. It is intended for IT leaders and others responsible for enterprise cloud strategy, cloud governance, and risk management decision making. This article summarizes key findings from the 2013 effort. It first discusses evolving concepts such as community clouds and cloud governance and identifies potential benefits and challenges DoD users face when migrating to the cloud environment. It then introduces a notional use case to help illustrate application of the Decision Method. Through the use case, the article explains how the needs and priorities of cloud users are considered when defining community-wide cloud policy. The same method can be used to negotiate policy updates as the environment and needs evolve.

The research described in this article was conducted in the 2012–2013 time frame while DoD CIO and the Defense Information Systems Agency (DISA) were in the process of operationalizing DoD's cloud strategy. DoD's cloud strategy continues to evolve and has recently changed significantly. Initially, the focus was largely on private and public clouds. Then in 2013 and early 2014, interest shifted to community clouds that DISA defined as cloud infrastructures that are provisioned for the exclusive use of the DoD and the U.S. federal government. DISA's Cloud Security Model and DoD's Risk Management Framework provide enterprise-level governance for security. But recently DoD announced that DISA is no longer the sole cloud service broker and the cloud security model has been revised as well.[5] However, even as emergent polices begin to abate differences in security practices, those differences related to consumer operations and priorities still persist. The Decision Method presented here can be used to address factors that a cloud consumer would consider when trying to determine the potential risk of operating in a given community cloud.

EVOLVING CONCEPTS

We used the National Institute of Standards and Technology (NIST) definition for community cloud for our research: "A community cloud serves a group of Cloud Consumers which has shared concerns such as mission objectives, security, privacy and compliance policy, rather than serving a single organization as does a private cloud. Similar to private clouds, a community cloud may be managed by the organization or by a third party, and may be implemented on customer premise or outsourced to a hosting company, Cloud Service Provider."[3] We define community cloud governance as a detailed strategy to translate mission objectives, rules, and decisions into precise and actionable policy statements that address the shared mission needs of the community. These statements may address access, orchestration, security zone, location constraints, as well as utilization rules and other thresholds having a potential impact on operations or capability performance. Effective community cloud governance ensures the execution of cloud resources in accordance with both users' operational needs and external constraints, addressing all phases of the system development life cycle while still maintaining the essential cloud characteristics: on-demand self-service, broad network access, resource pooling (in this case, between different organizations, also known as multi-tenancy), rapid elasticity, and measured service. (NIST defines cloud services by these five essential characteristics in *Special Publication 800-145, The NIST Definition of Cloud Computing*, which has become a de facto reference for cloud terminology.) The community cloud concept is one that continues to evolve and means different things to different business sectors. However, for this research, we adhere to the NIST standard definition. Likewise, the concept of community cloud governance is evolving as entities gain experience with cloud pilots and begin to understand how other cloud tenants' practices and operations may impact their own operations in a multi-tenant cloud environment. (Note that "there are varying degrees and definitions of multi-tenancy among cloud providers and many providers have the option of not sharing resources at an additional cost."[6])

The real challenge lies in understanding and identifying the policies that must change for the new dynamic and shared environment and then articulating and negotiating the new policy and appropriate security control capabilities, roles, and responsibilities. A new policy must address conditions at start-up as well as those that emerge as environments evolve. Currently, not all cloud providers offer what has come to be known as "enterprise-grade" management capabilities (e.g., visual policy editor, comprehensive application programming interfaces, integrated support for various cloud services, etc.). These capabilities support a broad vendor selection of cloud services for assignment of appropriate checks and balances across multiple stakeholders (e.g., IT, legal, and acquisition staff) or across multiple agencies, some of which cross national boundaries. To operationalize community clouds, the people, processes, and technical capability needs must be well understood to realize machine-speed cloud governance and security controls for this environment. (Machine speed here refers to the dynamic or near-real-time execution of policy changes by the implementation of business rules using an analytic engine. By

contrast, in a static environment where humans are involved in the transaction, a policy change may take hours or days to implement in the system.)

Cloud governance is an important aspect of the end-to-end planning, architecture design, deployment, and operations of a cloud project. Some key benefits and challenges that are unique to the on-demand, multi-tenant shared environment are shown in Table 1.

Table 1. The Case for Community Governance in the Cloud Environment

Benefits	Challenges
Increased IT efficiency and better cost management	Community-controlled governance activities need to be focused and prioritized
Improved agility, innovation, and resilience	Determining whether coordination and accountability should be centralized or distributed
Quicker operational deployment of new services	Identification of needed policy and supporting services changes
Increased productivity	Determination of appropriate framework for governance of community-wide concerns

It can be time consuming and costly to realize the appropriate balance of benefits while also addressing the unique challenges of this environment. But the adjudication of competing authorities and priorities can yield coordinated (and in the best-case scenario, integrated) operational objectives, policies, and strategies that can be codified for an efficient, effective implementation. Community-level engagement may be needed if implementation methods have potential to negatively impact the community or individual consumer operations. Consumers need to understand potential residual risks (defined by NIST as "the potential for the occurrence of an adverse event after adjusting for the impact of all in-place safeguards"[7]) to their operations and the system capabilities available to mitigate or minimize the risk to their critical operations.

COLLABORATIVE PLANNING ENVIRONMENT USE CASE

To explain the Decision Method, we consider the notional use case of a collaborative planning environment (CPE) that supports integrated operations. This case involves three organizations; each organization is a separate and distinct agency. Here, organization I is both a consumer and provider, while organizations II and III are consumers (see Ref. 8 for definitions of providers and consumers). Each has its own mission along with a shared mission that requires collaborative planning to develop response plans and courses of action. Figure 1 illustrates how, in addition to mission-related roles and responsibilities, each organization has responsibilities as a provider or consumer in the community cloud. The cloud provider and cloud consumer share the control

Step	Information Exchange	Speed
1	Organization I is identified as the lead actor for the planning of a given situation.	Human
2	Organization I analytics identify any existing planning materials that organization I can receive for this situation and pulls them via the CPE framework.	Machine
3	Organization I collaborates with organizations II and III to develop or modify the integrated response plan.	Human
4	Organization I coordinates plan execution with organizations II and III.	Both

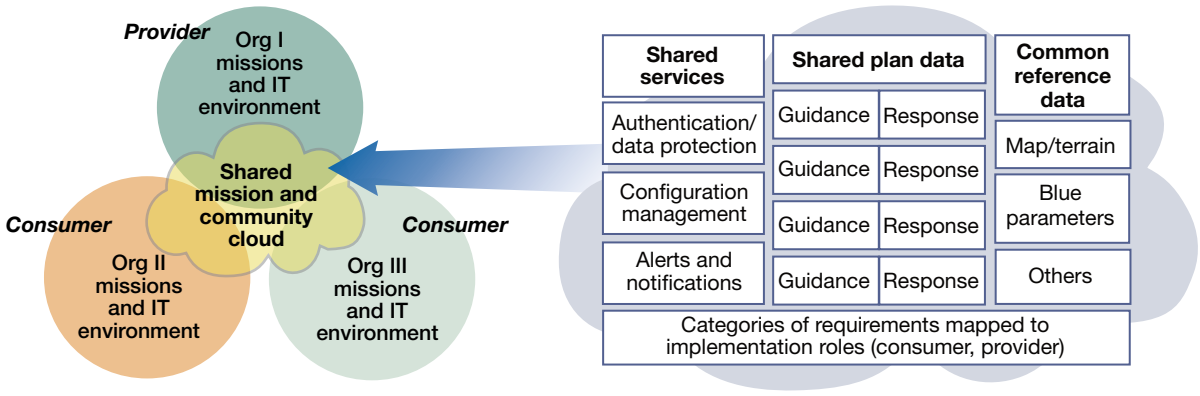

Figure 1. Use case information exchange and shared mission needs.

of resources in a cloud system. Different service models determine an organization's control over the computational resources and thus what can be done in a cloud. The planning function involves a number of information exchanges that are consumer or provider controlled.

The environment can be governed in various ways. The research hypothesis was that communities need to govern certain aspects of the cloud environment at the community level. The methodology discussion illustrates how this might be achieved for the use case considered.

DECISION METHODS

Decision methods are developed to address two major objectives. The first is to identify and assess common operational priorities with respect to the elements of the cloud's operational environment. The second is to quantify the degree of compatibility of cloud users based on operational norms. We will briefly discuss how each of the requirements is analyzed by the decision support analysis tools (DSATs), which leverage mathematical techniques from information theory to highlight elements of decisions while reducing the introduction of unintended biases.

Assessing Community Formations—Priority Compatibility

The initial step in deciding whether one or more organizations have similar needs with compatible objectives is to define the major factors that will be considered. For each of these factors, state tables will be built that comprise their structures. First, the relative importance of each factor (i.e., its weighting) is established for each participating organization; this may be accomplished by several means (see Ref. 9 for techniques to eliminate unintended bias in weighting).

An organization's factor weightings are a vector that is normalized so that the sum of the elements is unity. If the normalized vectors between two organizations are identical or closely similar, then the premise is that these organizations will be compatible in their demands from a community cloud and they may effectively negotiate for services and live happily in the operational environment.

To carry out objective comparisons of organization factor weighting vectors, we leverage the information theory concept of relative entropy, or the Kullback–Leibler distance.[4] This metric is a measure of the difference between two distributions. The vector of normalized factor weightings is a distribution, and if the distance between two organizations' vectors has a distance of zero or some small value, we would judge that the organizations are compatible in their views of what is important. Let p represent the normalized factor weighting vectors of an organization and q that of a different organization; the relative entropy between p and q, denoted as $D(p\|q)$, is defined as

$$D(p\|q) = \Sigma\, p(x)\, \log_2 (p(x)/q(x)),$$

where x ranges over the elements of the distributions. Because this expression is not symmetric, $D(p\|q) \neq D(q\|p)$, and, thus, the order in which the organizations are evaluated may confuse the analysis. To avoid this complication, we use instead the sum of two expressions, $D(p\|q) + D(q\|p)$, and denote it as $D(pq)$.

The value for relative entropy is always positive, but it is zero if and only if $p(x) = q(x)$. If $D(pq)$ is zero or very low then we may judge that the emphasis they put on decision factors is identical or very similar.

In this example, suppose we have five organizations that are considering working in the same cloud service environment. For each organization there are eight factors, (a) through (h) (the nature of the factors is identified in Table 2), that have various levels of importance, with the weightings shown in Table 2. The normalized weighting vectors are contained in the right part of the table and are computed by summing the weightings of each organization's column and dividing each element by the sum.

Table 2. Example: Five Organizations, Eight Factors

Factor Symbol	Factor Name	Organization					Normalized Weighting Vectors				
		I	II	III	IV	V	I	II	III	IV	V
(a)	Data impact level	4	3	7	1	8	0.133	0.071	0.125	0.036	0.200
(b)	Access control	1	4	8	2	8	0.033	0.095	0.143	0.071	0.200
(c)	Data, application sharing, integrated processes	3	2	3	3	8	0.100	0.048	0.054	0.107	0.200
(d)	External interfaces	5	8	8	8	5	0.167	0.190	0.143	0.286	0.125
(e)	Mission-based utilization	4	8	8	8	5	0.133	0.190	0.143	0.286	0.125
(f)	Criticality	5	6	7	3	3	0.167	0.143	0.125	0.107	0.075
(g)	Incident management and reporting	5	6	7	2	2	0.167	0.143	0.125	0.071	0.050
(h)	Compliance data visibility	3	5	8	1	1	0.100	0.119	0.143	0.036	0.025

Table 3. Relative Entropies of Organization Values

Pairing	D(pq)	Percentile	Fig. 2 Symbol
I–II	0.255	2	
I–III	0.335	4	
I–IV	0.738	23	
I–V	1.046	41	
II–III	0.125	0	
II–IV	0.458	8	
II–V	1.114	45	
III–IV	0.834	29	
III–V	0.796	26	
IV–V	1.099	44	

The results of carrying out the relative entropy computations are contained in Table 3 and illustrated in Fig. 2.

Other than perhaps a notion that organizations I, II, and III seem to be closer in their values than other subsets, without further context we may puzzle at the significance of these values. We can examine the distribution of possible values for this example by setting up a simple Monte Carlo analysis: consider the two distributions, p and q, and let the weightings of each element vary over a range of 0^+ to 1^- (we want to avoid division by zero in our computations, so avoid 0 and 1 for any given element); formulate the corresponding normalized vectors, compute D(pq) over many trials (say 100,000), and examine the cumulative distribution of the values for D(pq)—see Fig. 2—and learn what percentiles the D(pq) values represent.

Note that the 10th percentile cutoff is arbitrary and can be mandated by policy, but it does represent a 90% chance of having organizations forming compatible communities. The clustering of the D(pq) values for the I–II–III factor vectors falls below the 10th percentile of D(pq) values, indicating that these organizations share similar values with respect to the importance of factors. The II–IV organization pairing's value (black diamond in Fig. 2) also falls below the 10th percentile; however, organization IV's D(pq) values with respect to I and III (the purple diamond and the gray circle in Fig. 2) are above the 20th percentile, indicating an incompatibility of organization IV with these other two.

With this analysis, we determine the first major issue of identifying organizations with similar values with respect to the factors (termed the priority compatibility). Now we turn to the issue of comparing how well the governance in a cloud meets an organization's requirements (termed the policy compatibility).

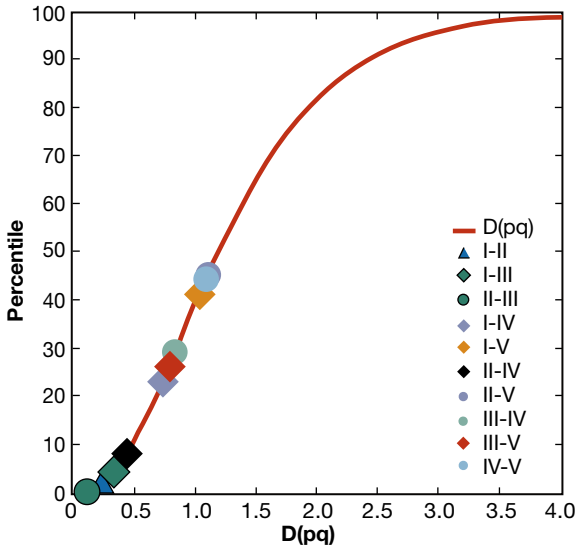

Figure 2. Cumulative distribution of D(pq) for two eight-element normalized vectors with example results.

Table 4. Factors and Their Elemental Yes/No Questions Used in This Example

Factor	Yes/No Questions ("Atoms")
(a) Data impact level	Q1: Is level 1? (Unclassified public; approved for public release)
	Q2: Level 2?
	Q3: Levels 3 to 5?
	Q4: Level 6?
	Q5: Level 7?
	Q6: Level 8?
	Q7: Level 9?
(b) Access control	Q8: Is control centralized?
	Q9: Decentralized (federated)?
	Q10: No access control?
(c) Data, application sharing, and integrated processes	Q11: Are data shared?
	Q12: Are applications shared?
	Q13: Are processes integrated?
(d) External interfaces	Q14: Are there no interfaces?
	Q15: Are there few interfaces?
	Q16: Are there more interfaces?
(e) Mission-based utilization	Q17: Can a surge be accommodated?
(f) Criticality	Q18: Is the criticality of the use negligible?
	Q19: Low?
	Q20: Medium?
	Q21: High?
	Q22: Very high?
(g) Incident management and reporting	Q23: Is everybody in the community notified (else only the two, CSP and tenant)?
(h) Compliance data visibility	Q24: Can members see other member's data?
	Q25: Are everybody's compliance data visible to everybody?
	Q26: Do members fence their data?

Table 5. State Table for Factor (c): Data, Application Sharing, and Integrated Processes

Level[a]	Logical Sentences in Level	Logical Sequences	MIE Weighting	Weighting Media
A	1	~x^~y^~z	1	1
B	1	x^~y^~z	3	3
C	1	~x^y^~z	5	5
D	1	x^y^~z	7	7
E	1	~x^y^z	9	9
F	1	x^~y^z	11	11
G	1	x^y^z	13	13

Factor (c) has three yes/no questions; $2^3 = 8$ logical sentences.
x = data are shared (yes/no); y = applications are shared (yes/no);
z = processes are integrated (yes/no).
[a] A to G = easiest to hardest requirement.

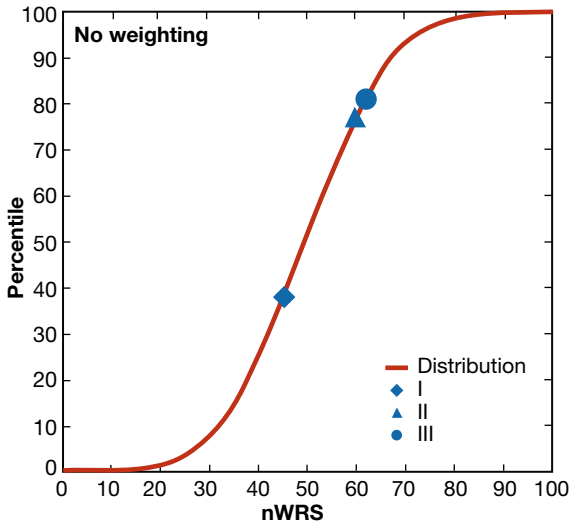

Figure 3. Cumulative distribution and illustration of requirements for organizations I, II, and III.

Quantifying Community Norms—Policy Compatibility

The factor details involved in our example, other than their names, heretofore have remained underdefined. Now we turn our attention to considering their specific content and illustrate a set of state tables built for evaluating a cloud policy. One feature of the technique is to reduce the input part of an evaluation to a series of binary (yes/no) questions, the answers to which determine a candidate's (cloud service's) value with respect to the whole universe (i.e., the allowable solution space) of allowable responses.

Table 4 introduces the factors that are considered in these examples. The specific factors used here are not intended to limit the scope of any future analysis and are described only to illustrate the technique to use when constructing a decision space. We keep the examples at a general level so as not to preclude limited distribution.

The 26 questions in this example are kept mostly plain just to illustrate the structure of state tables. For some of the factors, a yes answer to one of the questions precludes answering yes to any of the others in the factor group [e.g., for factor (a) there are only seven possible acceptable combinations out of the 2^7 combinations of yes/no answers that can be formed[10]]. For other factors, a richer set of elemental sentences can be constructed for the state table. Table 5 illustrates the state table for factor (c).

Each table is constructed from one or more yes/no questions (considered as logical atoms having one of the binary values of true or false). As part of the presentation format, the specific questions are identified below each table along with their representative symbols that appear in the logical sentences.

There are 2^n combinations of logical sentences formed from the n atoms. Depending on the nature of the system, some of these combinations may be impossible and are considered outside the solution space. For those feasible combinations, a sorting of the sentences, which now represents a gradient of requirements, is made from easiest to hardest. Note that some combinations may be judged to be equally desirable and a sorted into the same state.

Within a table, an assignment of weightings for each state is made using the maximum information entropy principle.[10] The factor weightings, conditioned so that each factor has an equal influence on a utility measure throughout the solution space, are shown in Table 6.

Table 6. Factor-Level Weightings for the Example

	Factor							
Level	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)
A	15,300	83,300	15,300	83,300	187,425	29,988	187,425	44,100
B	45,900	249,900	45,900	249,900	562,275	89,964	562,275	132,300
C	76,500	416,500	76,500	416,500		149,940		220,500
D	107,100		107,100			209,916		352,800
E	137,700		137,700			269,892		
F	168,300		168,300					
G	198,900		198,900					

For any given set of requirements among the eight factors, the utility measure (termed the weighted raw score, or WRS) is computed using the following equation:

$$\text{WRS} = \Sigma_{\text{all Factors}} \text{ (number of levels in factor} \times \text{factor level's weighting} \times \text{factor's weighting).}$$

All factor weightings are equal to 1 in this example (i.e., the factors are equally weighted). The WRS is normalized (nWRS) using the maximum and minimum values in the solution space:

$$\text{nWRS} = (\text{WRS} - \text{WRS}_{\text{minimum}})/(\text{WRS}_{\text{maximum}} - \text{WRS}_{\text{minimum}}).$$

From Table 6, there are 35,280 (i.e., $7 \times 3 \times 7 \times 3 \times 2 \times 5 \times 2 \times 4$) combinations for nWRS; not all values of these combinations are unique. The cumulative distribution of the scores is illustrated in Fig. 2, where the ordinate is the percentile and the abscissa is 100 nWRS.

An organization's requirements, captured as the answers to yes/no questions, are evaluated by computing the organization's corresponding nWRS. The values for organizations I, II, and III, the group that is determined to be a compatible community, are plotted on Fig. 3.

Figure 3 shows that organizations II and III have comparable demands on the CSP; organization I has a less intense set of requirements, although (from the priority compatibility analysis) it places similar values on the factors to those of the other two.

Summary of the Decision Methods

From this two-part examination of organizations' priorities and their individual requirements with respect to each factor, we may conclude that organizations II and III are fairly compatible and could operate in a shared community cloud. Although organization I has similar priorities, its demands within the community are different.

The foregoing analyses are embodied in a two-part DSAT, the community cloud governance (CCG) DSAT part I and the CCG DSAT part II, which are separate Excel workbooks that perform the detailed computations described above. In CCG DSAT part I, a user specifies the relative weighting of the factors; in part II, yes/no answers are provided to each of the 26 questions embodied in the state tables. The computations in part I give an assessment of the compatibility of the organizations with respect to factors; those of part II give the assessment of the requirements demand on the CSP and in relation to each other. These tools can be used immediately to analyze community cloud needs and can be augmented with updated factors and questions in future analyses.

UNDERSTANDING RESIDUAL RISK

As previously stated, the DSAT component of the Decision Method addresses two core issues: priority compatibility and policy compatibility. Evaluation of the first issue gives insight into whether the community has a chance for success as it forms and evolves. It provides a visual display of the areas of influence and their interdependencies: What is the current state? How can the current state be affected by operating variables (i.e., connectedness—physical and virtual), dependencies, and interdependencies? Reference 11 provides a detailed process for analysis of operational risk, based on an understanding of criticality and dependencies. The operational context for the methodology is similar conceptually in that it recognizes a greater need for understanding of interdependencies and connectedness within a community. The second issue indicates to a community member whether the cloud's policies to address critical factors are congruent with the specific requirements of the member. These analyses can be used to guide a potential member's negotiations with the other potential members in forming the community cloud. Modifications of the priorities of the factors and/or the specific policies chosen to address each factor may be part of that negotiation.

To facilitate effective collaboration in the CPE, agreements are needed on the rules of engagement. These rules are dependent on the identities, roles, objectives, constraints (e.g., policy, laws, financial), and claims of the parties involved, as well as their relative associations. Each consumer is identified by its profile and DSAT part I and part II representations. Comparison of these products can help determine how to most effectively design, implement, and manage the information infrastructure for the CPE and the agreements and enforcement mechanisms needed to ensure mutual cooperation, collaboration, and effective risk management of the environment. The degree of transparency needed for risk-based decisions is an example of a community-level governance decision.

For the CPE example, we concluded that organizations II and III are fairly compatible in sharing the same CPE. Organization I (which also serves as the provider of the cloud environment) has a measurable difference in its priorities and its requirements. Through these analyses, differences may be revealed as the parties negotiate the specific policies that the community would use to address each factor. Not every member of the community would achieve equal satisfaction. A member or group of members may decide to prioritize specific policies that improve their individual satisfaction, maybe to the detriment of the other members. Ultimately, a member must be satisfied with both the community agreement of defined factor priorities and specific policy implementations before deciding to

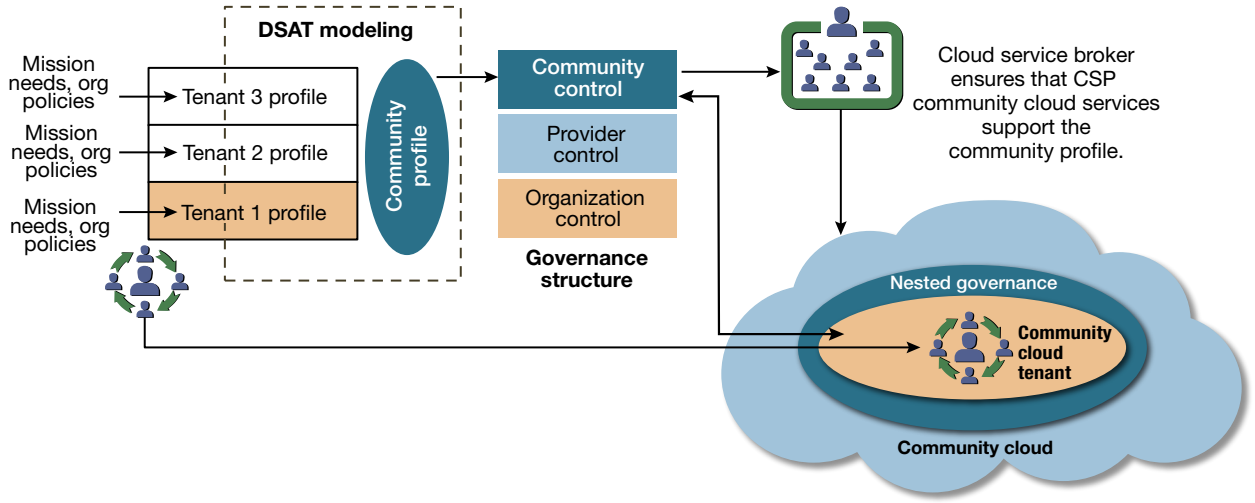

Figure 4. Multilevel governance structure for community clouds.

move to or remain in a cloud community. In cases in which an organization is mandated to operate in a given community cloud, the organization is at least better informed on how the cloud environment might impact its operations. Figure 4 shows a multilevel governance structure for community clouds. In this configuration, the cloud service broker may manage the degree of transparency provided for risk-based decisions. For other communities, direct control of risk-based decisions may be more appropriate.

CONCLUSIONS

The research described in this article led to development of a method an organization can use when deciding whether to join a community cloud. We identified the challenges that the Community Cloud Decision Method can address for community cloud decision makers and answered questions such as:

- How does an organization evaluate whether other potential consumers have common priorities that would make them compatible cloud co-tenants?
- How does a potential consumer determine what capabilities are needed to manage mission assurance concerns in a dynamic and shared environment?

Organizations migrate their business operations to community clouds for different reasons. An organization in need of cloud services may prefer a community cloud of like consumers. Their dilemma is how to determine that the other potential tenants in the community are in fact compatible. A different scenario is one in which an internal decision maker mandates that the organization move operations to a community cloud as a cost-savings measure or to support mission objectives (e.g., intelligence agencies required to share situational awareness information).

Consumers who elect to migrate to a community cloud might want the community to manage factors that are a high priority for most consumers. The assumption is that having the community make decisions about these factors helps the community's overall health status. The converse could be true. Consumers mandated to form a community cloud may differ significantly in critical factors. In this case, it might be prudent for the community to manage these factors with the utmost care to mitigate the risk for the collective. This is especially true for factors that significantly impact security or operations. This research offers a tool to help decision makers evaluate their options.

ACKNOWLEDGMENTS: The research leading to these results was the product of an independent research and development activity funded by the APL's National Security Analysis Department. In addition to the authors, other contributors to this research include Richard Bernstein, David Harper, and Bryan Canter.

REFERENCES

[1]Chief Information Officer, *Cloud Computing Strategy*, Department of Defense, Washington, DC, July 2012.
[2]Kunda, V., *25 Point Implementation Plan to Reform Federal Information Technology Management*, The White House, Washington, DC, December 9, 2010.
[3]Liu, F., Tong, J., Mao, J., Bohn, R., Messina, J., et al., *NIST Cloud Computing Reference Architecture*, Special Publication 500-292, National Institute of Standards and Technology, Gaithersburg, MD, September 2011.
[4]Cloud Security Alliance and ISACA, *2012 Cloud Market Maturity Study Results*, ISACA, Rolling Meadows, IL, 2012.
[5]Konkel, F., "Here's What the Rewrite of DoD's Cloud Strategy Will Look Like," *NEXTGOV Newsletter*, November 25, 2014.
[6]National Security Agency, *Cloud Computing Considerations*, October 2013, www.nsa.gov/ia/_files/factsheets/I43V_Slick_Sheets/Slicksheet_CloudSecurityConsiderations_Web.pdf.

[7]Wilson, M. (ed.), de Zafra, D. E., Pitcher, S., Tressler, J. D., and Ippolito, J. B., *Information Technology Security Training Requirements: A Role- and Performance-Based Model*, Special Publication 800-16, National Institute of Standards and Technology, Gaithersburg, MD, April 1998.
[8]Mell, P., and Grance, T., *The NIST Definition of Cloud Computing*, Special Publication 800-145, National Institute of Standards and Technology, Gaithersburg, MD, September 2011.
[9]Boyd, A., "New Rules Could Speed Up DoD Cloud Migration," *Federal Times*, April 15, 2015.
[10]*Risk Based Methodology for Verification, Validation, and Accreditation (VV&A), M&S Use Risk Methodology (MURM)*, NSAD-R-2011-011, The Johns Hopkins University Applied Physics Laboratory, Laurel, MD (April 2011), Appendix F2.
[11]Egli, D., "Risk Map," Chap. 3, *Facing the Storms: Operationalizing Preparedness and Critical Infrastructure Resilience*, D. Egli (ed.), The Johns Hopkins University Applied Physics Laboratory, Laurel, MD, 2013.

Monica M. Waters and **Christine O. Salamacha** were the principal investigators for this independent research and development project. Monica Waters, a member of the Senior Professional Staff at APL, provided overall direction of the research and has significant experience in information assurance and the application of advanced technology to military and intelligence problems. Christine Salamacha, a member of the Principal Professional Staff at APL, leveraged her extensive knowledge of DoD operations and the DoD cloud service broker initiative to contribute to development of the Decision Method. **Peter P. Pandolfini** is a member of the Principal Professional Staff at APL and an expert in modeling and simulation. He developed the DSAT, which has applications for many complex problem sets. All three authors work in APL's National Security Analysis Department. For further information on the work reported here, contact Monica Waters. Her e-mail address is monica.waters@jhuapl.edu.