Perspectives on Resilience in Cloud Computing: Review and Trends

Thomas Welsh, *Member, IEEE,* Elhadj Benkhelifa, *Member, IEEE,*
Cloud Computing and Applications Research Lab
School of Computing and Digital Technologies, Staffordshire University, Stoke on Trent, UK
(thoma.welsh; e.benkhelifa)@staffs.ac.uk

***Abstract*—The development of resilient distributed systems is seen as essential to maintaining stable business and state-run processes due to information systems now underpinning most aspects of society. Cloud computing is now one of the most pervasive usage paradigms and due its novelty, research surrounding its resilience is largely lacking and often varied in terms of developed solutions. Therefore this paper provides an up-to-date review of resilience work in cloud computing. This includes methods of measuring and evaluating resilience, solutions for enabling resilience and alternative architectures developed with a focus upon ensuring resilience from the ground up. Firstly, resilience is defined within the context of cloud computing in order to categorise the work appropriately.**

***Keywords*—Cloud Computing, Resilience , Alternative Architectures,**

I. Resilience Disciplines in Cloud and Conventional Computing

Resilience in the context of computer systems and networks, is known by a number of terms, often the exact description of which differs depending on the context and who is defining it. For some it is considered synonymous with, or a measure of fault-tolerance [1]. Some examples of the varied definitions follow. [2] gives two descriptions of resilience: ”the persistence of dependability when facing changes” and ”the persistence of service delivery that can justifiably be trusted, when facing changes”. Both definitions describe the persistence of dependability, although the second may be more open to interpretation. They are in fact describing a number of other desirable features of the system, as dependability contains a number of sub fields, and is also often considered a subset of trustworthiness. They define ”changes” as any ”failures, attacks or accidents” although arguably this may be considered somewhat ...

The author in [3] explains that this definition of resilience describes anything outside the system boundary, whereas dependability metrics often describe those within the boundary. A similar definition is given by the authors in [4], where they describe resilience as ”the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation”.

Faults can develop into system failures in the chain of *faults*, *errors* and *failures*; which describes the process in which *system failures* can occur. The measurement of services failures applies to the majority of the disciplines that will be discussed so it is important to understand before analysing any terms further. A failure, or system failure, occurs when a service deviates from its correct state and no longer provides the correct service; a result of not meeting it current specification or the specification being inadequately laid out [5] [4].

An error occurs within a system when it is observed to be different from the correct system state. An error is caused by a fault, which may be intentional or non-intentional, and external or internal. Hence the chain process which may lead to a service failure whereby a fault may cause an error, and an error may lead to failure. Other definitions are also given, such as [6] where ”Resilience is the capacity of critical services to adapt in order to provide their functionalities in cases of undesired events compromising parts of the system.” This definition mostly focuses on modelling the capacity of the system and thus could be seen as a dependability focused measurement as it concentrates on internal resources. It is similar to the definition given by [7] which focuses on resilience from a business/organisation perspective. ”Resilience refers to the capacity of human beings/system/organization to survive and thrive in the face of adversity”...”it is a property that is closely associated with the capacity to avoid, contain and mitigate accidents.” Again, this definition focusing upon measuring the capacity of the system which is likely concerned with the redundancy inherent within the system, redundancy is often a main characteristic of fault-tolerance. However this fails to take into account deliberate acts which are a component of trustworthiness and also the measurement of characteristics which cannot also be quantified by capacity.

In [8] the authors simply describe resilience as ”the percentage of lost traffic upon failures” another highly quantifiable measure which fails to take into account other factors such as those related to trustworthiness. [4] provide a survey of resilience disciplines in which they describe resilience as an umbrella term. According to these authors, resilience is composed of two major sub fields, with each one containing a number of other disciplines. These two sub-fields are grouped into those with relating to *trustworthiness* and *challenge tolerance*. Dependability, (along with security) is grouped into trustworthiness disciplines. Grouping of these attributes is similar to the previous definitions in that the attributes are arranged depending upon their placement relative to their interaction with the system. However, this definition more clearly defines the cross overs between sub-fields and the grouping into trustworthiness and challenge tolerance is more

logical and defined. A classification of these disciplines is given in Table 1.

The variety of definitions is perhaps due to the variety of different fields and disciplines in which resilience could be measured. Databases, data centres, computer networks , power transmission networks and business management are just some examples of this. As each one will have different inherent characteristics, a number of different models for measuring resilience appear. Additionally, specific use cases may omit certain characteristics due to their lesser relevance.

Cloud computing is a subset of computing and therefore inherits these disciplines, and with it the methods of measuring and achieving them. However, a major way in which cloud environments differ from the traditional paradigm, within the context of the aforementioned disciplines, is through its dynamic nature and its use-case based variations, driven by service requirements. In conventional computing systems and networks, when attempting to adjust or measure the resilience, it is likely that the internals of the system will not change too much and therefore the resilience requirements and features will also not vary considerably. Cloud environments, however, do vary due to their constant and sometimes poor determinism. Which is directly affected by their constant varying use-cases. It may be difficult to define exactly what resilience means in terms of computer systems, in general and for cloud computing, in particular. For the purposes of this work, it is considered that resilience is the master set of all aforementioned disciplines. Where each discipline is either a quantification, or feature enabled through a particular implementation, which helps describe the overall resilience of the system. This is adhering to the definitions as given in [4]. Within the context of cloud resilience, it is necessary to understand this chain as *system-failures* are the challenge in which the system must be resilient against.

II. State-of-the-Art in Cloud Resilience

The previous sections highlighted the complexity and contentious nature surrounding the variety of resilient characteristics and disciplines. In short, resilience may be considered context specific, in that the varying use cases should directly drive the requirements. The combination of resilience within cloud computing environments thus complicates the situation further due to the open-ended uses available for cloud environments. Therefore in a general sense, as inherited from traditional computing, resiliency of a cloud environment could be the quality of its trusted service delivery. Where quality varies depending upon the exact service constraints, where trusted is an assumed requirement that must be verifiable and the delivery is a boolean value. The ratio of successful, verifiable delivery is therefore a measure of the resilience of the cloud system.

Like all computing environments, cloud environments consist of multiple layers. What may distinguish the layered nature of cloud environments from the layered aspect of conventional computing environments, is that the layers within cloud may have different stake holders and are managed by different actors. This is considered different to the delivery models of cloud environments (IaaS/PaaS/SaaS) which are also layered but are concerned with pricing and service as opposed to the application, control or measurement of resilience.

Therefore due to this layered nature of cloud architectures, the following survey of cloud resilience categorises work according to the layer in which the work focuses upon. This is important to the direction of the study as although works may sometimes employ similar techniques, the focus of the resilience has a direct effect on specific layers and may be employed different stake holders. Therefore this categorisation occurs due to the service-oriented nature of cloud environments but with a focus upon the management of resilience within that layer. Currently there is no process for the investigation of resilience in the cloud so this layer based investigation was decided in order to categorise and compare the current literature. The layers were developed through an iterative process during the literature survey. The concept of the model is based upon the general layered model of cloud service delivery but due to the variations in resilience is adapted. It is similar to the classification of attacks and defences proposed by [9] which categorises the threats and defences according to their intended target. In this variation focused upon resilience, the defences are defined by their intended target but also within the context of their stake-holders. As the resilience of each layer may be directly altered by different stake-holders the work is grouped as such.

The proposed layers being as follows:

- **VM - Instance Resilience** - the work in this layer describes resilience techniques which focus upon providing resilience to an individual instance such as a virtual machine or container.
- **Multi-VM - Platform Resilience** - these works are concerned with cloud services which are composed of multiple instances and the way in which resilience may be managed between these instances.
- **Data-Centre - Management Layer resilience** these works are concerned with data centre management techniques which may provide resilience via the interconnection of multiple physical machines.
- **Data-Centre - Networking and Physical Layer Resilience** The works in this layer are concerned with the physical layer (e.g. hardware) and networking within the datacentre.

Most of reported research within the area of cloud resilience focused upon specific layers of the cloud environment, as shown in Tables I-IV. Some other work have modelled different factors which affect resilience within cloud infrastructures; while fewer research has been reported where proposal for new additions or or complete changes in cloud architecture were advocated. The latter research forms part of the focus of our research, where the intention is to shed light on some promising trends in cloud computing resilience with completely alternative architectures. This is further discussed in the next section

TABLE I. SUMMARY OF RESILIENCE DISCIPLINES

Discipline	Requirements	Description	Type	Superset
Survivability	Diversity and Redundancy	Facing component and system failures	Feature	Challenge Tolerance
Fault Tolerance	Redundancy	Recover from random faults	Feature	Challenge Tolerance, Survivability
Traffic Tolerance	Redundancy, Scalability, Bandwidth	Survive varied increases in traffic	Feature	Challenge Tolerance
Disruption Tolerance	Routing Techniques	Disruption of communication paths	Feature	Challenge Tolerance
Dependability	reliability, availability, maintainability, safety	Quantification of Service Delivery	Quantification	Trustworthiness
Reliability	?? -	Continuity of service	Quantification	Trustworthiness - Dependability
Maintainability	Ability to change	The ability to undergo system changes	Feature	Trustworthiness, Dependability
Safety	?? -	Dependability during critical failures	Feature	Trustworthiness, Dependability
Security	Availability, integrity, confidentiality, nonrepudability, auditability, authorisability, authenticity	Protection against unauthorised change or access according to policy	System Property	Trustworthiness
availability	??	Readiness for service	Quantification	Trustworthiness, Security, Dependability
Integrity	Cryptography and verification	"absence of improper system alterations"	Feature	Trustworthiness, Security, Dependability
Confidentiality	Cryptography	The ability to keep data and actions secret	Feature	Trustworthiness, Security, Dependability
Nonrepudability	Cryptography	Ensuring a user cannot deny an action or reciept	Feature	Trustworthiness, Security, Dependability
Auditabiltiy	Accurate Logging	Ability to keep accurate logs	Feature	Trustworthiness, Security, Dependability
Authorisability	Polcies, Credential Management	Granting and denying permissions to resources according to policies	Feature	Trustworthiness, Security, Dependability
Authenticity	Cryptograph	Ensuring the integrity of message, sender, content and meta data (time etc.)	Feature	Trustworthiness, Security, Dependability
Performance - QoS				

TABLE II. INSTANCE LAYER RESILIENCE

Work	Method	Feature
[10]	Reactive	error ranking and appropriate technique
[11]	Reactive	introspection
[12]	Proactive	high diversity for replica storage
[13]	Proactive	memory stored backups

TABLE III. PLATFORM LAYER RESILIENCE

Work	Method	Feature
[14]	Service Composition	Graph based, interdepdency
[15]	Service Composition	Agent-based
[16]	Low Level Diversity	data-centre
[17]	Quality Adjustment	brownout
[18]	Diversity	Replicas
	Diversity	Structure height

TABLE IV. INFRASTRUCTURE MANAGEMENT RESILIENCE

Work	Method	Feature
[19]	Organisation	VM Scheduling for resilience
[20]	Disaster Recovery	Storage
[21]	Reactive reset upon fault	Hypervisor
[22]	Proactive Diversity	Hypervisor
[23]	QoR Evaluation	Proactive service provisioning

TABLE V. INFRASTRUCTURE NETWORKING RESILIENCE

Work	Method	Feature
[24]	Diversity	Geo-distribution
[25]	VM to PHY Mapping	Design
[26]	VM to PHY Mapping	Backup links
[27]	Redundancy	VNet backup links

III. ALTERNATIVE ARCHITECTURES FOR CLOUD COMPUTING RESILIENCE

Some work will choose to encourage a conventionally different cloud architecture in order to provide increased resilience (Fig. 4). Although being an alternative to the infrastructure layer discussed in the previous section, the authors present SlapOS [28], choosing to provide a purely distributed cloud architecture where the issue of single point of failure is remedied through distributed the cloud resources over multiple PCs within homes, as opposed to within Data Centres. Whilst this might obviously bring forth issues regarding bandwidth,

TABLE VI. RESILIENCE THROUGH ALTERNATIVE ARCHITECTURES

Work	Method	Feature
[28]	Decentralisation	SOHO Device architecture
[29]	Decentralisation	Community Cloud
[29]	Decentralisation	Leader Selection Optimisation
[30]	Diversity and Redundancy	Multi-layer service aware manner
[31]	Diversity	Constant evolution
[32]	Diversity	Execution and I/O redirection
[33]	Redundancy	Mature components with internal redundancy
[34]	Redundancy	Resource prediciton
[35]	Diversity and Intellignece	Novel Service provisioning

capacity and latency the benefits of reducing single point of failure are unparalleled during decentralisation, particular for safety-critical events such as those during disaster events.

Resilience within SlapOS is again the focus in [29]. The authors highlight the lack of resiliency within a conventional IaaS cloud as motivation for their work and the development of SlapOS. The authors reiterate the lack of resilience within current cloud architecture's due to the centralised data centre model. They refer to the concept of community cloud, whereby the cloud is collaboratively built from personal devices. The main current issues are summarised as:

- Migrating from commodity cloud to resilient, secure and dependable clouds
- Promoting diverse and open ecosystems
- Building a coherent, modular and reusable architecture

When considering a distributed cloud, the leader selection problem is highlighted [29] (the process of selecting the next master node after loss of the current). Whilst this a valid issue, it perhaps signifies the need to me to purely distributed cloud architectures. Further issues relating to resilience of their application include: implementation an accurate failure detection methods, and methods of replicating the master database prior to handover to another master node. An interesting point made by the authors is that using the SlapOS architecture, the conventional delivery models of infrastructure, platform and service delivery become obsolete. Finally, the authors explain that an implementation of hierarchical masters (such as with DNS) will be implemented for increased resilience. Whilst the architecture and delivery model is certainly interesting with this work there are issues directly relating to resilience concerning master node hierarchies which undoubtedly cause problems. A decentralised system such as this is not as resilient as one which is purely distributed.

Community cloud based resilience is also discussed in [36] which promotes the model as an enhancer for organisational resilience. As with the work of SlapOS, the authors highlight the ownership and location issues of current cloud models being unsuitable for providing resilience. A point of interest by the author is that for natural disasters centralised disaster recovery tends to become too late and too excessive, the author argues that disaster recovery must be conducted by the community level. The breakdown of communication networks is cited as a key issue here, where the more effective communication was developed by the decentralised communities. The author explains that community cloud models enable all the benefits of public cloud offerings whilst enabling greater control. Issues surround community clouds, such as malicious users, are said to be mitigated through user vetting, a process which may not always be practical or effective.

An architecture known as DefCloud [30] attempts to provide greater resilience through increasing diversity and redundancy within all layers of the cloud architecture. The architecture is also flexible, in that it allows resilience to be adjusted in a "service-aware manner". This might be argued to be similar in concept to the usability vs security trade-off. Such a feature is likely necessary for a cloud platform which accommodates a wide spectrum of use-cases. The first key point in designing the infrastructure, as argued by the authors, is the removal of monoculture which, for example, enables malware and attacks to propagate effectively through only needing to attack one type of hardware architecture or software application. This reduction of monoculture is then applied to all layers of the cloud infrastructure. Firstly is the *infrastructure layer diversity*. This is encompasses *data-centre diversity* and *cloud diversity*. Where data-centre diversity is best considered as sub-trees of features where similar trees are not selected in tandem in order to maximise diversity. For example similar trees will utilise the same network vendor hardware or operating systems. Whilst diversity provides resilience against security related failures, it does not protect against failures due direct physical data centre attacks, e.g. natural disasters or military attacks (such as an EMP). In order to mitigate these issues, the architecture then applies cloud diversity through distributing the cloud over multiple geo-locations, using varying ISPs.

After the infrastructure layer, the DefCloud then assures resilience through *Process-level Program diversity*, where the diversity focuses upon distribution via space and time. Where the spatial diversity is concerned with distribution of differing versions across the cloud whilst the temporal diversity is concerned with varying application configurations over time. In short, application diversity ensures individual binaries are diverse meaning that an attack on one application binary will not apply to another. Whilst this has profound consequences on the current state of 0-day exploits, it creates a large number of issues for the software development process. Whilst the architecture uncountably covers resilience in the cloud through heavy adaptions of the conventional architecture, the system lacks real implementation or simulation and thus its resilience is yet to be determined. For one, the complexity of the system is clearly greater and therefore the number of attack points rises also.

Similar diversity may be seen in the MEERKATS system [31] which as a fully novel architecture for cloud security, which focuses on a security mission critical cloud. The system constantly evolves across all aspects, reducing monoculture and increasing diversity. One component of the system, DREME [32], is concerned with execution diversity of replicas and provides a framework for I/O redirection.

IBM present a somewhat novel architecture name SCE+ [33] which is built from the ground up to be highly resilient. The authors make the distinction between typical cloud architectures employed by Amazon and Google by explaining that they are constructed from "redundant, inexpensive, expendable building

blocks" whereas the IBM SCE+ employs "high-end building blocks with significant internal redundancy and an established track record of very high MTBF for every element." In short, it would appear that the contrast is in SCE+ employing mature and extremely resilient fewer components with conventional architectures employing many less mature components and relying upon replication/redundancy.

The architecture applies resiliency to differing layers within the cloud. The physical layer is designed so as to avoid single point of failure, through division of resources and replication in separate geo-locations with a backup dark-fibre link. Software resilience is then considered from multiple aspects. Components are deployed in redundant pairs and constant "health-checks" are in place to monitor correct functioning. In addition, redundancy of data and regular backups ensures resiliency within the data layer. The authors go on to explain that standardisation of hardware within the system components aids the resilience, however this is contentious, as diversity within hardware is surely a necessity for resiliency. They also cite virtualisation as an enabling factor of the resilience, however this a typically a component of cloud infrastructures anyway and therefore offers the environment no additional advantage. Overall the architecture offers a variety of additional components for resilience although some are questionable such as the physical distance between components as well as the added complexity within the system.

Moving away from a traditional cloud architecture, improving the resilience of Hybrid Mobile Clouds is the focus in [34]. Mobile clouds require greater resilience than a static system due to the dynamic network characteristics. The proposed architecture is interesting due to its flexibility in running on a variety of device types, essentially ignoring the underlying hardware. The resilience requirements are also aided through a resource prediction mechanism and an early failure detection mechanism to facilitate handover of vital services. The system proves successful, although performance is still dependent upon the quantity of fixed nodes within the cloud, making the system not purely mobile. However, overall it exhibits a good example of how cloud systems can be built upon non-deterministic environments.

An architecture based on biologically inspired processes which allows tunable redundancy at multiple cloud levels, known as BioRAC is presented in [35]. One layer of the architecture involves division of components into "cells" which allows dynamic real-time configuration and combine together to form an "organism" which then is then applied to a particular goal. In an additional layer, the system provides high levels of diversity through varying execution and finally it provides intelligent algorithms for collaborative thread alert and detection. Although lacking an implementation or proper evaluation, the architecture is interesting in providing a system designed with resilience from the ground up with novel components, as opposed to those adapted on top of conventional systems. However the system is undoubtedly complex due its multiple layers which has an adverse effect upon its complexity and thus its resilience. With lacking any implementation it is difficult to assess its resilience although the techniques and concept are certainly of merit.

IV. Analysis: Research Gaps and Future Directions

After an analysis of the previous work in Cloud computing resilience, a number of research gaps can be highlighted for future research and directions. This work can span multiple levels. But specifically, is concerned with resilience in cloud environments, which are disparate from the traditional cloud architectures. As mentioned previously, traditional cloud environments are essentially resilient by nature, which is mostly due to the high redundancy involved but also due to the ability for a user to customise a cloud service or application with their own features and then enable this resilience according to their own SLA with the provider.

What seems to be largely missing from the literature is research which focuses upon cloud resilience in more constrained and less deterministic environments. Such work is highly relevant due to the nature of emerging cloud disciplines e.g. mobile cloud, fog computing and edge computing. Such constrained environments have less ability to fall back upon redundancy in order to provide their resilience and must rely on more complex methods. They might also employ a variety of diversity related techniques due to disparate hardware involved. A key factor which is deemed relevant to the growing field of cloud computing, which is not fully investigated, is how does the effect of resilience upon one layer, affect the resilience of another layer? For example, if resilience is enabled by a user in the platform/service oriented layers but the underlying physical layer has low resilience, is it still possible to increase resilience in this manner? Such a topic is highly relevant to the way in which cloud architectures are evolving to more mobile and less deterministic networks and away from highly deterministic data centre environment. Another area which is not touched upon sufficiently is the ability to dynamically adjust within constrained environments, whilst some work within engineering has focused upon applying dynamic algorithms to graph analysis and optimisation, little work has been conducted which leveraged this for the cloud environments. Once more, this has particular relevance for mobile environments due to the constrained resources available for optimisations, machine learning has seen some efforts in this area but further work can involve an evaluation and comparison of different algorithms for both traditional and mobile cloud environments.

Another area which could be expanded upon is the theoretical nature of enabling resilience within the context of various constraints. This has particular relevance to cloud SLA but also to constrained models. It concerns the analysis of requirements to enable the degree of resilience for the service. As resilience is a scale as opposed to a binary value, such a model could aid the construction of a service within its given constraints across all cloud service models.

In terms of physical layer resilience, the exact effect upon resilience in the cloud with different levels of diverse hardware has seen minimal work. Therefore, future research directions in this area could see the exact effect of diversification of hardware resources upon the resilience of a system could be investigated. Barriers to this research mostly involve cost and time; as the necessary hardware, proprietary licenses and

TABLE VII. BIO-INSPIRED DOMAINS MEETING FUTURE DEMANDS FOR UNCONVENTIONAL CLOUD COMPUTING RESILIENCE

Bio-Inspired Technique	Information Capacity	Network Adaptability	Network Convergence	Structural Redundancy	Structural Diversity
Epigenetic	High	Medium	Finalised	High	Low
Phylogenetic	Medium	Medium	Dynamic	High	Low
Ontogenetic	Low	Medium	Dynamic	High	Medium
Swarm - Societal	Medium	High	Dynamic	High	Low

practical work involved in evaluating these scenarios ensures it is difficult to implement. However, simulations may enabled a realisation of evaluating this approach.

Despite being used as inspiration for a variety of optimisation and resilience methods throughout engineering and computing, minimal work can be seen within the area of cloud resilience which leverages bio-inspired work. Biological systems have been shown to have innate resilience, which may seem particularly relevant due to the high-complexity of these systems, which is similar to the application of resilience in cloud systems. Therefore, this lack of work highlights a huge potential area for further application of biological systems to cloud computing environments.

Undoubtedly, biological systems provide resilience as in the definition adopted above. Nature has found ways for life to persist, guaranteeing the delivery of a ”service” (or life) in the face of many changes. However, to fit in with the proposed model, which treats resilience as an umbrella term for a variety of sub-fields, nature also provides this resilience across varied and diverse use-cases. The huge number of different scales and types of life such as animals as an entity, swarm-based systems such as ants, and more complex systems (such as the earth ecosystem), which consist of a variety of sub-organisms; has shown that a number of scales of resilience may be achieved through differing characteristics, depending on the use-case. Table 2 present and map the main bio-inspired dimensions with these characteristics.

Given the inherent resilience of biological systems presented in this section, this analysis seeks to understand which particular techniques they employ and how they are useful to resilience in computing. Therefore, an approach is taken which compares two key aspects of resilience: redundancy and diversity. In addition to the characteristics the techniques employ to optimise or converge upon a solution. A highlighted characteristic within literature, which enables their resilience, is their ability to adapt through learning, some being able to continuously adapt (swarm techniques), whilst others will eventually converge on a given solution (Genetics). The ability to converge may allow a system to become optimal at a particular solution, but leave it liable to attack. In contrast, a system which continues to adapt may never be able to provide appropriate functionality for one solution. Another characteristic is the quantity of information each system contains at both an atomic and network scale layer. For example, epigenetic nodes hold very little information and are largely concerned with functional/processing characteristics, such as with Neural Networks and Immune systems' simplicity at each node. In contrast, the nodes within an ontogenetic system will hold considerably more information, processing information locally with each node's distinct functionality. The inherent distributed nature of all biological systems ensures they are all the ability to apply scalable levels of redundancy rather effectively. However, diversity will vary, ontogenetic systems for example, provide considerably greater diversity within their atomic units and system than those of epigenetic systems which are considerably more similar to one another. Although the ability to create further diversity from the atomic unit should also be of note. For example, while the ontogenetic systems may develop into complex systems with considerable diversity than that of the epigenetic, these systems develop from the phylogenetic processes and if flaws exist from this process, they can propagate through the entire system, making it weaker therefore. Whilst biological systems can be seen to enable resilience, and optimisation, of different levels within computing; they are applied in different ways. Swarm-based systems, AIS and ANNs all apply a minimalistic algorithm, which was modelled from a biological process. Others, such as computer viruses take a broader concept and apply it to computing systems. Few works within literature are seen to apply multiple layers of a biological system to an architecture. This may be for a number of reasons, such as a lack of seen necessity, i.e. if only one aspect needs optimising within a conventional system then why replace all of it? It might also be due to the complexities of biological systems causing problems when attempting to model it. If it is difficult for scientists to understand the interactions between these layers, then it is likely too difficult for engineers to map them. Therefore, a point which should be considered during this application of biological techniques to man-made engineering; focuses upon applying models of complimentary layers of a biological system, in tandem. In order to create a bio-inspired system which is more considered to be like an ALife than just a single algorithm or application.

V. CONCLUSIONS

Resilience in computing systems is covered by a number of definitions, though it is mostly considered as the superset of a large number of sub-disciplines, which provide differing characteristics and functionalities to enable the given service. The three components of diversity, redundancy and intelligent organisation are recurring themes within this review and as such are considered key enablers of resilience. However due to the service driven nature of cloud environments the enabling methods of resilience vary considerably depending on the use-case. As such in order to provide a comprehensive resilience method it may be necessary to explore those which cover all layers which were previously discussed in this paper. Of particular note are those architectures which are unconventional, as traditional architectures have been shown to be somewhat resilient but only through expensive methods such

as redundancy. Therefore alternative solutions may be the most optimal method of providing holistic resilience solutions to the wide and varied use-cases prevalent within cloud computing environments.

REFERENCES

[1] W. Najjar and J.-L. Gaudiot, "Network resilience: a measure of network fault tolerance," *Computers, IEEE Transactions on*, vol. 39, no. 2, pp. 174–181, Feb 1990.

[2] J.-C. Laprie, "Resilience for the scalability of dependability," in *Network Computing and Applications, Fourth IEEE International Symposium on*, July 2005, pp. 5–6.

[3] K. Trivedi, D. S. Kim, and R. Ghosh, "Resilience in computer systems and networks," in *Computer-Aided Design - Digest of Technical Papers, 2009. ICCAD 2009. IEEE/ACM International Conference on*, Nov 2009, pp. 74–77.

[4] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Comput. Netw.*, vol. 54, no. 8, pp. 1245–1265, Jun. 2010.

[5] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *Dependable and Secure Computing, IEEE Transactions on*, vol. 1, no. 1, pp. 11–33, Jan 2004.

[6] C. Queiroz, S. Garg, and Z. Tari, "A probabilistic model for quantifying the resilience of networked systems," *IBM Journal of Research and Development*, vol. 57, no. 5, pp. 3:1–3:9, Sept 2013.

[7] N. Abdullah, N. Md Noor, and E. Ibrahim, "Resilient organization: Modelling the capacity for resilience," in *Research and Innovation in Information Systems (ICRIIS), 2013 International Conference on*, Nov 2013, pp. 319–324.

[8] G. Liu and C. Ji, "Scalability of network-failure resilience: Analysis using multi-layer probabilistic graphical models," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 1, pp. 319–331, Feb 2009.

[9] M. A. Khan, "A survey of security issues for cloud computing," *Journal of Network and Computer Applications*, vol. 71, pp. 11 – 29, 2016.

[10] T. Nguyen, J.-A. Desideri, and L. Trifan, "Applications resilience on clouds," in *High Performance Computing and Simulation (HPCS), 2012 International Conference on*, July 2012, pp. 60–66.

[11] F. Lombardi, R. Di Pietro, and C. Soriente, "Crew: Cloud resilience for windows guests through monitored virtualization," in *Reliable Distributed Systems, 2010 29th IEEE Symposium on*, Oct 2010, pp. 338–342.

[12] H. Reiser and R. Kapitza, "Hypervisor-based efficient proactive recovery," in *Reliable Distributed Systems, 2007. SRDS 2007. 26th IEEE International Symposium on*, Oct 2007, pp. 83–92.

[13] R. Jhawar and V. Piuri, "Fault tolerance and resilience in cloud computing environments," *Computer and Information Security Handbook,*, pp. 125–141, 2013.

[14] M. Mihailescu, A. Rodriguez, C. Amza, D. Palcikovs, G. Iszlai, A. Trossman, and J. Ng, "Enhancing application robustness in cloud data centers," in *Proceedings of the 2011 Conference of the Center for Advanced Studies on Collaborative Research*, ser. CASCON '11. Riverton, NJ, USA: IBM Corp., 2011, pp. 133–147.

[15] V. S. Sharma and A. Santharam, "Implementing a resilient application architecture for state management on a paas cloud," in *Proceedings of the 2013 IEEE International Conference on Cloud Computing Technology and Science - Volume 01*, ser. CLOUDCOM '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 142–147.

[16] P. Verissimo, A. Bessani, and M. Pasin, "The tclouds architecture: Open and resilient cloud-of-clouds computing," in *Dependable Systems and Networks Workshops (DSN-W), 2012 IEEE/IFIP 42nd International Conference on*, June 2012, pp. 1–6.

[17] C. Klein, M. Maggio, K.-E. Årzén, and F. Hernández-Rodriguez, "Brownout: Building more robust cloud applications," 2014.

[18] M. Guo and P. Bhattacharya, "Diverse virtual replicas for improving intrusion tolerance in cloud," in *Proceedings of the 9th Annual Cyber and Information Security Research Conference*, ser. CISR '14. New York, NY, USA: ACM, 2014, pp. 41–44. [Online]. Available: http://doi.acm.org/10.1145/2602087.2602116

[19] J. Cartlidge and I. Sriram, "Modelling resilience in cloud-scale data centres," *CoRR*, vol. abs/1106.5457, 2011. [Online]. Available: http://arxiv.org/abs/1106.5457

[20] V. Jaiswal, A. Sen, and A. Verma, "Integrated resiliency planning in storage clouds," *Network and Service Management, IEEE Transactions on*, vol. 11, no. 1, pp. 3–14, March 2014.

[21] A. Binun, M. Bloch, S. Dolev, M. Kahil, B. Menuhin, R. Yagel, T. Coupaye, M. Lacoste, and A. Wailly, "Self-stabilizing virtual machine hypervisor architecture for resilient cloud," in *Services (SERVICES), 2014 IEEE World Congress on*, June 2014, pp. 200–207.

[22] M. Kanter and S. Taylor, "Diversity in cloud systems through runtime and compile-time relocation," in *Technologies for Homeland Security (HST), 2013 IEEE International Conference on*, Nov 2013, pp. 396–402.

[23] B. Sousa, K. Pentikousis, and M. Curado, "Methodical: Towards the next generation of multihomed applications," *Computer Networks*, vol. 65, pp. 21 – 40, 2014.

[24] V. Westmark, "A definition for information system survivability," in *System Sciences, 2004. Proceedings of the 37th Annual Hawaii International Conference on*, Jan 2004, pp. 10 pp.–.

[25] M. Bui, B. Jaumard, and C. Develder, "Anycast end-to-end resilience for cloud services over virtual optical networks," in *Transparent Optical Networks (ICTON), 2013 15th International Conference on*, June 2013, pp. 1–7.

[26] M. Bui, T. Wang, B. Jaumard, D. Medhi, and C. Develder, "Time-varying resilient virtual network mapping for multi-location cloud data centers," in *Transparent Optical Networks (ICTON), 2014 16th International Conference on*, July 2014, pp. 1–8.

[27] I. Barla, K. Hoffmann, M. Hoffmann, D. Schupke, and G. Carle, "Shared protection in virtual networks," in *Communications Workshops (ICC), 2013 IEEE International Conference on*, June 2013, pp. 240–245.

[28] G. Suciu, C. Cernat, G. Todoran, V. Suciu, V. Poenaru, T. Militaru, and S. Halunga, "A solution for implementing resilience in open source cloud platforms," in *Communications (COMM), 2012 9th International Conference on*, June 2012, pp. 335–338.

[29] R. Courteaud, Y. Xu, and C. Cerin, "Practical solutions for resilience in slapos," in *Cloud Computing Technology and Science (CloudCom), 2012 IEEE 4th International Conference on*, Dec 2012, pp. 488–495.

[30] J. Sterbenz and P. Kulkarni, "Diverse infrastructure and architecture for datacenter and cloud resilience," in *Computer Communications and Networks (ICCCN), 2013 22nd International Conference on*, July 2013, pp. 1–7.

[31] A. Keromytis, R. Geambasu, S. Sethumadhavan, S. Stolfo, J. Yang, A. Benameur, M. Dacier, M. Elder, D. Kienzle, and A. Stavrou, "The meerkats cloud security architecture," in *Distributed Computing Systems Workshops (ICDCSW), 2012 32nd International Conference on*, June 2012, pp. 446–450.

[32] A. Benameur, N. Evans, and M. Elder, "Cloud resiliency and security via diversified replica execution and monitoring," in *Resilient Control Systems (ISRCS), 2013 6th International Symposium on*, Aug 2013, pp. 150–155.

[33] V. Salapura, R. Harper, and M. Viswanathan, "Resilient cloud computing," *IBM Journal of Research and Development*, vol. 57, no. 5, pp. 10:1–10:12, Sept 2013.

[34] A. Khalifa, M. Azab, and M. Eltoweissy, "Resilient hybrid mobile adhoc cloud over collaborating heterogeneous nodes," in *Collaborative*

Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on, Oct 2014, pp. 134–143.

[35] S. Hariri, M. Eltoweissy, and Y. Al-Nashif, "Biorac: Biologically inspired resilient autonomic cloud," in *Proceedings of the Seventh Annual Workshop on Cyber Security and Information Intelligence Research*, ser. CSIIRW '11. New York, NY, USA: ACM, 2011, pp. 80:1–80:1. [Online]. Available: http://doi.acm.org/10.1145/2179298.2179389

[36] G. Garlick, "Improving resilience with community cloud computing," in *Availability, Reliability and Security (ARES), 2011 Sixth International Conference on*, Aug 2011, pp. 650–655.