

Diagnostic Work in Cloud Computing: Discussion Forums, Community and Troubleshooting

John Rooksby

School of Computer Science,
University of St Andrews, UK.
jrn@st-andrews.ac.uk

Ali Khajeh-Hosseini

Cloud Computing Co-Laboratory
University of St Andrews, UK.
akh@cs.st-andrews.ac.uk

ABSTRACT

As systems scale, systems management often becomes partially reliant on web forums and other social media. This paper examines the use of web forums for diagnostic work in cloud computing. We argue that forums are not simply used to communicate information but that (with users attempting to negotiate and manage the attention of providers, forming coalitions, criticizing others, and framing problems in particular ways) forums are socially organised, value laden venues for information. We conclude that providers should focus not on improving communication, but more broadly on managing community.

Author Keywords

Infrastructure as a Service (IaaS), Conversation Analysis.

ACM Classification Keywords

K.4.3 Organizational Impacts: CSCW.

INTRODUCTION

The cloud computing paradigm sees computational resources (software, data storage, servers, etc.) accessed as a service rather than acquired as a product. IaaS (Infrastructure as a Service) is one model of cloud computing [1]. IaaS providers enable more or less anyone with a credit card to build and run a system on servers owned by and located with the provider. The idea is that if large providers supply computing services over the Internet on a commercial, pay-as-you-go basis, to a massive, self-service user-base, users will benefit from substantial cost savings, improved uptime, and extreme flexibility.

Using an IaaS provider involves developing and operating a system across organisational boundaries. Using a larger provider means the user may be just one of many thousands of others, which has implications for the way support can be provided. One convenient technology for providing

support in this situation is the web forum. Forums provide a familiar, many-to-many line of communication, and simultaneously build a temporal, public archive of problems and solutions. This paper addresses how web forums are being used for identifying failures and planning remedial actions, for what Büscher et al [2] term “diagnostic work”. The paper concludes that in order to improve user support, large providers need to focus not on “communication” with users but more broadly on how they encourage and manage the cooperative work of its “community”.

DISCUSSION FORUMS IN CLOUD COMPUTING

Most IaaS providers maintain a discussion forum as a part of their support services. Smaller providers can rely on more direct contact if they wish, but the larger ones cannot. As the market in this area expands, reliance on this relatively cheap form of support is likely to increase.

This paper focuses on forums run by AWS¹ (Amazon Web Services), which is one of the larger and more mature IaaS (Infrastructure as a Service) providers. AWS provides free “basic support”, giving access to discussion forums, to technical FAQs, and a dashboard detailing service availability. Paid “premium support” gives access to one-to-one web-based and telephone-based support. We have found premium support does not replace the forums, they are a first point of call, a hub for linking to other information, and enable user-user interaction. Discussion forums seem a convenient technology, both for provider and users. AWS currently provides twenty-one individual discussion forums. The two most popular forums, for Amazon EC2 (Elastic Compute Cloud) and Amazon S3 (Simple Storage Service), have tens of thousands of messages and millions of views.

The general design of the AWS forums is much like that of any other professional forum. Messages are organized into threads in order of posting. Their functionality, along with the ‘house rules’, both enable and constrain users in what they can do. However this does not in itself determine how people use the forums, it merely sets the stage for socially organized interaction.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW 2012, February 11–15, 2012, Seattle, Washington.

Copyright 2011 ACM XXX-X-XXXXX-XXX-X/XX/XX...\$5.00.

¹<http://aws.amazon.com/> (Amazon, AWS, Amazon Web Services, Amazon S3 and Amazon EC2 are trademarks).

1 **User1: tons of 500 error messages** *Posted: 1:17 PM*
2 I am trying to upload a bundle and I keep getting 500
3 error messages like this one:
4 Failed to upload "/mnt/image.part.03":
5 Server.InternalError(500): We encountered an internal
6 error. Please try again.
7 ERROR: Failed to upload "/mnt/image.part.03":
8 Server.InternalError(500): We encountered an internal
9 error. Please try again.
10 Also I was trying to use S3Fox to delete things from my
11 bucket and I get the same 500 error messages. This all
12 usually works perfectly fine for me.
13 **User2: Re: tons of 500 error messages** *Posted: 2:28 PM*
14 Same here.
15 I've been unable to test my new image for the last 4 hours
16 because of the same error.
17 This problem started yesterday for me, when I had to re-
18 run the upload over 15 times before it would complete.
19 Well, it's Friday and there's no cloud. I might as well go
20 to the beach early :)
21 **User3: Re: tons of 500 error messages** *Posted: 2:46 PM*
22 Having the same problem here. I expected to find more
23 people complaining. It started this morning for me.
24 Simple gets and deletes started throwing 500s almost
25 consistently. Yesterday was fine.

26 **John@AWS Re: tons of 500 error messages**
27 Posted: 3:35PM
28 If you get a 500 response from Amazon S3, retry your
29 request using exponential backoff between retries. With
30 the EC2 command line tools, you can use the "--retry"
31 argument to have the tools retry your request up to 5
32 times.
33 You can find more information about dealing with
34 Amazon S3 errors in the documentation found here:
35 <http://docs.amazonwebservices.com/AmazonS3/latest/ErrorBestPractices.html>
36
37 If you continue to have problems after employing these
38 measures, please provide the request ID (if available) and
39 the bucket name for the failing request.
40 **User2 Re: tons of 500 error messages** *Posted: 4:58 PM*
41 Hi John,
42 this was using the api tool ec2-upload-bundle.
43 It seems it does not follow the ErrorBestPractices then.
44 I worked around the problem by no longer bundling
45 amis.
46 Instead I started applying patches stored on s3, on top of a
47 reference ami.
48 This has a faster turnaround time. Uses less traffic and is
49 so far keeps me more productive than having my hands
50 tied with bundle upload error.

Figure 1: Example thread (Anonymised and Formatted)

MUNDANE AND CATASTROPHIC FAILURE

In this paper we will focus on a mundane failure. Reason [4] argues that mundane failures are normal in complex systems, and that catastrophic failures usually involve unfortunate combinations and timings of familiar, ordinarily mundane factors. Catastrophic failures are certainly important to understand, but we argue 1) mundane failures are common and so important to understand in their own right, and 2) catastrophic failures are best understood with reference to mundane failures.

Catastrophic failure (of a kind) can and does happen in cloud computing. On the 21st April 2011, AWS suffered an outage in one of its data centers, taking around two days to resolve. In turn, a large number of systems built on AWS failed or were adversely affected. According to AWS² the failure was initiated by human error within the data center. AWS admitted that problems for users were then compounded by a lack of communication; users found it difficult to identify what was happening and what actions, if any, they should take. Studying the forums will not help in preventing future outages but does give insight into how information is shared and how lessons are learned by users. In response to the April 2011 outage, AWS has stated it is

working on “improving communication”. This is important, but through the lens of mundane failure we suggest there is a wider issue. How can the way knowledge is constructed and shared by providers and users of the forum be improved? We suggest the answer lies not just in “communication” but also more broadly in “community”.

ANALYSIS OF A MUNDANE FAILURE

In this paper we draw from Antaki et al [3] to examine how posts to the AWS forums are socially organised. We will do so by focusing on the on-stage interactions; we will not try to explain what is going on behind the scenes at the provider’s or users’ workplaces. We focus on-stage because the forums are boundary objects situated between people who are unlikely to know each other and potentially have little shared interest except for the service itself. Antaki et al [3] have argued that the coherence and orderliness, and ultimately the meaningfulness and usefulness of a forum, must be achieved through the interactions that take place in writing on the forum itself, through how participants “make their messages achieve recognizable social and personal objectives while attending to the discursive perils” (such as rejection, criticism or silence).

Following Antaki et al [3], we will draw from ethnomethodology and conversation analysis to discuss a single thread from the AWS S3 forum (figure 1). We are

²Our comments are based on the official postmortem from Amazon <http://aws.amazon.com/message/65648/>

describing interactional not cumulative phenomena and so have chosen an example that gives relatively clear, concise demonstrations of phenomena we believe are salient. The example is slightly longer than many threads and relatively jargon-free, but is otherwise *in our experience* a run-of-the-mill thread. We have been visiting cloud forums (run by AWS and other providers) over a two-year period. Our analysis has involved three multi-disciplinary data sessions. The forums (including the example thread) are publically accessible and our findings reproducible.

Opening Message

The opening message in the example (lines 1-12) is entitled “*tons of 500 error messages*”. This title serves as succinct problem statement. It is not designed to lure readers in (as might be a headline or advertisement) but written to make the message content unsurprising, to only attract the attention and time of those directly interested in the message content.

The message body describes how an error has started to occur (lines 2-3), illustrates this with copy and pastes of error output (lines 4-9), and is then backed up by describing another situation in which the same error has also started occurring (lines 10-12). This opening message, as is usual in any forum [3], is more than a factual problem statement. The problem is framed in three ways; it is said to be happening: a) repeatedly - “*I keep getting*” (line 2); b) in more than one context - “*Also I was trying to use S3Fox... and I get the same...*” (lines 10-11); and c) for something that has worked previously - “*this all usually works perfectly fine for me*” (line 12). Contextualizing the problem in these ways serves to justify its inclusion in the forum; it: a) gives a demonstration of competence; b) demonstrates the user has tried several possibilities before seeking help; and c) indicates the problem is most likely with AWS and not with the user’s implementation. Ultimately the message is, in Antaki’s terms [3] “accountable”; it is trying to solicit a particular account from the provider (an answer), gives the resources the author thinks are necessary for coming to that answer, and contextualizes the problem so as to avoid criticism and rejection.

Confirmations (and Challenges)

The second and third messages (lines 13-25) in the example thread are not answers but “me too” confirmations. Both messages state the same problem is happening and, as with the opening message, add contextual information that points to a particular diagnosis. Both refer to the fact that the problem is repeatedly encountered after multiple attempts: a) “*I had to re-run ... over 15 times*” (lines 17-18); and b) “*It [is] ... throwing 500s almost consistently*” (lines 23-25). Both also refer to the problem as being something that started at such-and-such a time: a) “*The problem started yesterday*” (line 17); b) “*It started this morning*” (line 23). Trying to identify consistencies is a method of externalising the problem, to locate it with the provider. The role of these confirmation messages is not to sympathise with the

first author, but to help solicit a response from the provider. Confirmation messages serve to bring attention to a problem: they a) show it is not an individual, one-off problem; and b) bump the message to the top of the forum.

There is a hint of emotion in the second message. The comment about going “*to the beach early :)*” (line 20) is good humoured, playing down the urgency of the problem. We’ve noticed that when there are many confirmations the opposite often happens, they often become angrier in tone and start magnifying the importance of a problem.

Interestingly, the third message states “*I expected to find more people complaining*” (lines 22-23). This reveals how confirmation messages can serve to reinforce a user’s suspicion that a problem lies with the provider. Because there are fewer than expected, the author now seems to doubt his/her initial suspicion.

Sometimes a message may challenge rather than confirm a previous one. These may directly dismiss a problem, or attempt to reframe it. Challenges serve as reprimands; they are a discursive peril, as discussed by Antaki et al [3].

Solution and Workaround

The example contains two answers. The first is from the provider. This answer begins by mentioning good behaviour in retrying requests: “*exponential backoff*” (line 29). This point is unsolicited; the provider is taking the opportunity to mention a good practice, one that helps them avoid getting swamped with retries. Next, the message points to documentation, leading readers away from the forum. The documentation itself actually contains very little extra information. The link serves not to provide extra details but as a strong indication that this problem is for the user to work out away from the forum. Finally, the alternative is to post information about a “*request ID ... and bucket name*” (lines 38-39). This recognizes the possibility of a problem with the provider’s service, but casts it as an individual rather than systemic issue. This answer closes down possibilities for appropriate responses. The user should a) go away and read the documentation, and b) failing that, give specific information to be used for a specific diagnostic act by the provider. User1 did their best to be a good customer, they pitched their problem in a respectful, meaningful way but the provider needs them to be a good customer on their particular terms: they want someone who checks the documents and then provides the information with which AWS can do the basic (and likely most effective) check.

The final response (to date) rejects the provider’s answer and gives a workaround. User generated solutions and workarounds are common on these forums, but are often presented as a dis-preferred type of answer. User solutions contain not just information but typically a) criticize a perceived failure of the provider to follow the line of reasoning the user(s) have been taking in troubleshooting and then present a workaround; or b) where a workaround

is offered in advance of an ‘official’ reply, it will be presented in a tentative manner, not closing down the possibility of an official response. As a rejection, the final message also touches on some of the emotional labour of diagnostic work: it employs an argumentative, almost angry tone.

FORUMS, COMMUNICATION AND COMMUNITY

In response to a catastrophic failure AWS have stated they will work to “improve communication”. We believe the way this catastrophe was handled should be put in the context of how mundane, everyday failures are handled. This way the wider problem of fostering and managing *community* in cloud computing becomes visible.

Community is a word already used by AWS: their website links to the forums with the words “tap into the breadth of existing AWS community knowledge” and “engage with and learn from the AWS community.” This study shows the term “community” is far from ironic. The term, however, does gloss over socially organized, interactional phenomena. Following Antaki et al [3], we have used a single example to illustrate how:

- Posts are matter of fact, but not statements of pure fact. They are written in what Antaki et al [3] call an “accountable” way (they are written to encourage certain next turns, while limiting the scope for others).
- Locating a problem (with provider or user) and deciding who holds responsibility for diagnosing and solving it can be negotiable and collaborative.
- Knowledge accumulates on the forum in a value-laden way. Hyperlinks are provided both as answers and for moving people away from discussion, and workarounds are presented as tentative.
- A ‘good customer’ for the provider is not just a polite one, but one who is able, without prompting and within a short message, to provide the information they need for carrying out their routine diagnostic procedures.
- Answers are oriented to a general audience (for example giving general points about good practice).
- Messages in the forums treat the attentions of the provider and other users as an extremely limited resource (correctly or otherwise).
- Interaction on the forum can draw upon and display emotion.
- The default assumption of the provider is that the problem lies with the user. Users form coalitions to dismiss this assumption (rightly or wrongly).

Büscher et al [2] point out that diagnostic work usually relies not on specialist technologies, but common, everyday ones. Because these technologies are relatively mature in design terms, the work done across them can often be overlooked. Discussion forums are one such technology.

There may be limited design challenges, but the work done across them needs to be better managed. This paper makes a step towards that by unpacking aspects of this work.

SOCIOTECHNICAL APPROACHES TO FAILURE

The field of cloud computing has focused extensively on design for failure, but too often from an overly technical perspective. We argue that preventing, mitigating and recovering from failure is a socio-technical not technical problem. The April 2011 outage underscored this for us. The outage was initiated by human error in a data center, and compounded for users by a lack of communication. Subsequently, there was a debate on the forums and in the technical media about responsibility; some argued it was Amazon’s fault that the systems built on AWS failed, others argued developers should have built their systems across multiple data centers. Irrespectively of where fault lies, an important point is that there are human, social and organizational issues at play. Dependability in cloud computing is a socio-technical, not purely technical problem. Addressing diagnostic work on the forums is just one small step into a much wider set of issues to which sociotechnical research can make useful contributions.

Some of the issues we have covered in this paper are familiar to CSCW. In particular there are some parallels here with studies of the use of forums in distributed software development (although in the case of cloud computing, the forum users are not working together on a project but are more self-interested, and relationships between users are likely to be more transient). Our ultimate intention with this work is not to describe some novel practice but to bring the knowledge and methods of CSCW to the development of cloud computing.

CONCLUSION

Diagnostic work across discussion forums relies not just on provider-user communication but on cooperative work between the provider and a community of users. If cloud computing is to become more reliant on discussion forums, we believe the social organization of this community needs to be understood and managed.

REFERENCES

1. Badger L, Grance T, Patt-Corner R, Voas J (2011) Cloud Computing Synopsis and Recommendations. National Institute of Standards and Technology, Special Publication 800-146.
2. Büscher M, O’Neill J, Rooksby J (2009) Designing for diagnosing. *Journal of Computer Supported Cooperative Work*, 18(2-3), 109-128.
3. Antaki, C., Ardévol, E., Núñez, F., and Vayreda, A. (2005) "For she who knows who she is:" Managing accountability in online forum messages. *Journal of Computer-Mediated Communication*, 11(1), article 6.
4. Reason, J. (2008) *The Human Contribution*. Farnham: Ashgate.