NETEVENTS

APAC PRESS & ANALYST SUMMIT

Final

Conference Debate VII:

From Millions to Billions of End Points: Stress-Testing the Cloud and the Internet of Things

Chaired by: Anshul Gupta

Research Director, Gartner

Panellists:	
Derrick Loi	Senior Director, DC Solution and Services, Orange Cloud for Business, Asia Pacific
Amit Sinha Roy	VP Marketing & Strategy, Tata Communications
Naveen Bhat	Vice President & General Manager Asia-Pacific, Ixia
Ashwin Jaiswal	Head - IT Business Consulting & Practice, Reliance Communications

My panellists are here: Derrick Loi, Senior Director, DC Solution and Services, Orange Cloud for Business, Asia Pacific; Amit Sinha Roy, VP Marketing & Strategy, Tata Communications; Naveen Bhat, Vice President & General Manager Asia-Pacific, Ixia; and Ashwin Jaiswal, Head - IT Business Consulting & Practice, Reliance Communications. Thank you so much for joining me today.

All right. So morning everyone. My name is Anshul Gupta. I work with Gartner and it's my pleasure to be here, hosting this session with my panellists. Today we are going to talk about stress-testing IoT nodes and the cloud. And I just want us all to imagine a scenario where it's a Sunday afternoon. It's raining outside; really gentle rain, tiny water droplets coming down your window, on the trees and so forth. It's really gentle; you can imagine that setting. You're sitting on your sofa or couch watching a game or playing -- or watching a movie, playing a game, or maybe it's a good time to take a nap. So these individual water droplets are very gentle in isolation, but these water droplets aggregate and turn into streams of water. The cities in the world have to plan for water run-offs. They have to build the infrastructure, drainage,

sewer system and so forth. Even the best-designed infrastructure in the world to handle water run-offs can be overwhelmed by floods of water, either because of too much of rain coming in at the same time, or maybe because of the infrastructure being overwhelmed.

So you might be thinking this is really an inappropriate analogy to use in this technology session, but I believe that there is a connection because what we are talking about is Internet of Things. So there are going to be sensor data; in isolation that sensor data is going to be very small, but that sensor data aggregates and turns into streams of data. And that streams of data can quickly overwhelm our system, infrastructure, applications, storage and so forth.

What we are talking about here is 21 billion connected things by 2020. And we are talking about [$1.9 trillion] of added economic value by 2020. But here is the real thing that really blows my mind, that data will be doubling every two years. We are talking about 33 zettabytes of data by 2020. So let me explain a bit here; it's mega, giga, tera, peta, exa and zetta. That's a lot of zeros in there. So we are going to be in a very complex situation when we talk about Internet of Things. We are already seeing that our communication service providers are experiencing network outages. In 2014 March EE in the UK experienced a nation-wide outage. Their customers could not use voice, SMS, data for several hours. We had a similar scenario in 2014 April in US where Verizon suffered outage of their emergency services in California. Customers in nine counties of California could not use emergency services during that period of time. So these network outages are not always because of the technical issues. Sometime there is also a human angle involved into it. Think of a scenario where the network performance is affected because in the neighbourhood there is NFL's Super Bowl game going on; there are thousands of people who are streaming, sharing high mega-pixel pictures on the social media, and bringing down or affecting the network performance.

So this brings forth the importance of testing. There are various angles to look at when we are talking about testing in an IoT scenario. If we look at the connectivity, where we talk about connecting things to the network infrastructure, there will be performance of the network itself or the devices itself; there will be security aspect, that is very, very important, how to really secure the data, maintaining the privacy; functionality, dealing with the user interface, embedded softwares; compatibility with the Internet of Things we are going to see a number of communication protocols, number of devices coming from different vendors. So there is going to be a lot of challenges on that side.

So today I am going to focus more on connectivity and on the performance side. So I would like to ask my first question to Amit here. Considering that soon there are going to be billions of devices connected on to the Internet, and that means we will also deal with the risk what if those devices start misbehaving. And since you are also in an early stage of setting up an IoT network in India, so I just wanted to understand your views on the importance of testing and also what solution you have really come up with.

Amit Sinha Roy

Thanks for that. So when we talk about testing, we have to look at it from a holistic perspective which you have put up, right? So there is end-user, right, who is involved and the testing is actually to provide the user experience to that end user in the most seamless manner. So if I divide it between the various aspects of testing, there is firstly the part that comes obviously is the last mile which is what we're doing with the LoRa LPWAN-based network, right? And if there is a deployment that is being done with the LPWAN LoRa technologies, there is guidelines, there are testing methodologies that is provided as part of the LoRa Alliance, very specific, well documented, that is used for testing various aspects in terms of the range, in terms of the frequency, in terms of various other operating parameters which is there, which obviously we're doing to get the last mile, in the most robust way, rolled out.

But from there it comes on to our network, right, and then from our network obviously it goes on to the Internet and to the other application and providers who are there. So that's in terms of the transmission of the information. But it doesn't -- the whole stress-testing just doesn't stop there, because it's -- actually when you look at the whole value chain, it starts with the embedded device which is there or, if it is a smart device, that itself, what is it doing? Then if it is a scenario which yesterday we alluded to, supposing somebody has a smart fridge and you have auto replenishment, right, and say it is for cheese. The cheese manufacturer then has to make sure that the sensor that they're embedding into that box has the right characteristics. And the fridge manufacturer has to ensure that it can read that. Then there is an application provider who has to write the application to ensure that they are able to capture that information. And then, ultimately, if it is auto-replenishment, there is an e-commerce or there is a supplier involved as well.

So in between that there is the network which passes the data on, right. But what if there is a problem with the tag and it just starts spewing out data? It could be almost like a DDoS tag because it keeps sending out a bunch of unwanted packets, and which has happened, because if you go back, I think, a couple of years ago, there was a huge issue on one of the networks and it actually turned out to be a refrigerator that was spewing out information which was not required and just overloading the network, right? So from that perspective then, the whole play of stress-testing is not just a single point and it goes across various stakeholders to give that user experience at the end of the day which would be seamless in terms of what they actually want out of it.

But from our side, certainly in the rollout we're doing the last mile testing of the LoRa LPWAN network that we've rolled out across three cities, making sure that it meets all the parameters as per the guidelines, as per the framework. And then of course, from our own network experience, we route 24 percent of the world's global Internet traffic, so there is a good team that we have who is managing that. And I'm sure they'll be able to manage this as well.

Anshul Gupta

Yes. So I think that -- you have made a very good point about connected refrigerators and the other items really tweeting and really affecting the networks at times. And I would like to extend this question to Naveen. You deal with lot of OEMs, testing their products. And then you also deal with the CSPs and the other enterprises as well. So given -- I mean, I understand that when OEMs are coming up with a product, they will be doing the extensive testing. But then there is another round of testing which is really needed when you really put up into a network, because there could be different configurations where those products are being deployed. So how do you really help communication service providers, doing those kind of testing, where the scenarios could really be different, but the product be the same.

Naveen Bhat

Hi. Yes, this is Naveen. To answer your question Anshul, I think you need to start thinking about, first, the granular level of device testing, where each individual device is tested against one another and how well it talks. Then you start going into groups of devices and how well they talk to one another and how well they communicate back. Then you start talking about thousands of devices and then millions of devices, right. And part of what Ixia does is provide the capability for a manufacturer to test a simulated set of traffic coming from thousands or millions of devices. And, unless somebody can say, with confidence, that I've tested it out, okay, if a thousand devices hit my network, if a million devices hit my network, my network will stay alive, right, if they cannot say that with confidence, then we have a problem in the actual rollout right?

And let me point out one other thing; in your chart there you talked about connectivity and performance, but there's one other element too which is latency. So as you start adding more and more devices into the IoT network, which is right now at about 5 billion, expected to go into 50 billion very soon, then the issue of latency becomes really important because there are certain types of devices which are more in the entertainment space, and refrigerator and TV space, but there are other devices that are much more life critical and life threatening. And in which case if you any latency, then you've got a serious problem. So testing it out before deployment is essential and that's what we provide tools for by which people can test thousands of devices or millions of devices.

Anshul Gupta

Yes. So that's another good point about the latency and that's where I really think about -- thinking about the performance of the network. And, in an IoT scenario, there are, at times, scenarios where the network really has to be very, very agile; where you need to do the fast deployment at times and when there some scenarios you really need to deal with it. So that really brings forth the importance of adding intelligence into the network through the SDN or self organizing networks and all. And these sensors or the devices at times are also very hardware constrained. They don't have enough processing power; the batteries. So how do you really optimize

that network for that battery performance? I think that's another important parameter, so I would like to ask that question to Ashwin. Maybe you can share your views on that.

Ashwin Jaiswal

Yes, hi. This is Ashwin Jaiswal from Reliance Communications, Head of IT Business Consulting. Yes, I think there are multiple -- as Naveen also pointed out that latency becomes a crucial aspect of the overall testing, there are also some aspects which are, for example, there are multiple protocols under development. There are protocols which are already rolled out like, for example, LoRa and SigFox etcetera. But there is something which is coming up in Release 13 also. Now all of these devices, as we were discussing yesterday, that why don't somebody just shut off the [IPv4] version and just move to 6 because if you have finite numbers of it, then we can straight away move. But then the fact is that they have to be backward compatible.

Now, similarly, in this scenario also, if you look at it, until and unless you have backward compatibility with multiple protocols, finally any solution will -- as a telecom operator, it has to work with multiple devices, multiple versions, multiple platforms and multiple OS of them. Now, they have to be -- they have to work cohesively at one single point in time, because, let's say, there's a HP device of medical equipment which is working with, let's say, Siemens. Siemens may have their own platforms. This will have their own platforms, but finally it comes to a telecom network. Let's say that a hospital has solution from, let's say, it's Reliance. It has to work with all of these devices, all of these platforms. That testing, those protocols have to be very, very important.

There is another aspect to it also. For example, there are -- these equipments, whom we are talking about, they are connected things. If you take a manufacturing scenario, there will be big, large machines who will be placed very deeply inside. Similarly in hospitals, let's say, CT scan, the machines are kept generally in the basement area or inside area. How will they communicate finally with the networks. Now, there will be multiple hops; either they will use internal WiFi, then probably connect. So testing scenarios, as Naveen mentioned, is not going to be so easy. It's going to be -- when we talk about case-to-case basis and when we're talking about 50 billion devices, how are you going to test them? It's not that they're going to test 50 billion devices, but the fact is that those scenarios have to be -- somebody has to analyze those scenarios; somebody has to put thought processes behind it.

It's all evolution right now. And until and unless we reach to that point, it is going to take certain amount of time for everybody to start accepting it. So currently it's going to be business -- case-to-case basis. But obviously, as we go along and we are -- and we start taking up the pace, I think the scenarios will be very, very different. And it's going to be mind-boggling, as rightly shown in the picture, the drops becomes a flood. And this is how it's going to be a huge flood. And although it may not sound very positive to talk about, but the fact is that most of the operators will face certain kind of outages at certain point of time before we actually land up, because no testing in the world has ever proven that that can be 100 percent successful. It has to fail

somewhere and failure is something that will teach all of us. And I think all of us will pass through that -- those failures and we'll keep learning.

Anshul Gupta

Thank you. So I would like to ask my next question to Derek, because he is here and he has been handling the cloud side of the business. And when we are talking about IoT, the cloud really plays a very important role in designing the whole IoT solution. So I would like to have your views; how do you scale the cloud for IoT implementation?

Derrick Loi

Orange Business Services prides itself on the fact that we have tailored our cloud solutions to suit different verticals as well as different use cases. And one of the use cases that we have recently developed our cloud portfolio for is actually big data. So I will use the analogy of an actual customer to answer the question. I can't mention the name, but this customer is from the hotel industry, so they have hotel chains around the world. They wanted to engage Orange to deliver a globally available and consistent managed big data-ready platform, so that they can then define the right policies/rule-sets to collect and analyze data from various sensors and their booking system to make sense of their customers' profile, their preference, the usage of each and every hotel, the facilities within hotels and, obviously, how they should then, correspondingly, manage and adjust their resource and headcounts to support that.

So the first step that we took was, how do we make our network and our data centre/cloud services analytics-ready. Now, as the diagram earlier showed, it is no joke; when you have a huge amount of data that is collected by multiple sensors and multiple devices, flooding the network and flooding the data centre, so that's when a raindrop becomes a flood. And therefore, it was important, especially for this hotel chain, to first carry out certain amount of pre-processing and filtering at what we call edge analytics, in order to really extract and optimize the data before they are sent back to the central data centre for further processing and analytics. Now, this pre-processing and filtering of edge analytics that's done at all the hotel branches, this is where we discovered the first pain-point. A lot of these hotel branches were ill-equipped to handle that sort of edge analytics and the associated volume of data. So this was where we implemented the first wave of our project - to upgrade all the infrastructure at the branch office, what we call branch optimization.

We actually worked with the customers to collate and consolidate some of their existing Web servers, proxy servers, database, mail servers, even the routers and the WiFi access point, into a single appliances or dual appliances for high availability. And these appliances are now capable of actually hosting and running the IoT applications that are needed to actually collect and pre-process and filter the data

that's collected by all the sensors and all the devices running at each hotel branch. So that was the first step that we took.

Now, beyond that, what was also important was a lot of these hotels demand that their operations are uninterrupted, especially when there's a WAN failure. So this is where, again with this branch optimization exercise, not only do we now enhance their on-premise compute, storage capability for IoT, we also now enable the ability for the branch hotels to not suffer any disruption. So the collection of the data, the ongoing transactions for their day-to-day business will continue, regardless of whether there's a network outage. And the moment that network connectivity is resumed, that's when all the data is synced and replicated back into the central data centre. So that was the first pain-point and the first step that we took to address it.

Now, the second step was definitely towards providing a analytics-ready platform that allows them to host, run and set up instances of Hadoop, Spark on demand. As you know, there are essentially three types of IoT applications and their transactions. There's the metrics and measures, there's the actual transaction itself and there's the analytics data. Especially for the first and the third type of data, they're highly bursty, very transient in nature; whereas for the second, which is the transaction data, these are of the most value, but yet they are also the most prevalent and voluminous, they demand the most in terms of the data consumed.

So each of these data type in IoT would require different treatment. This is where we worked with the customer, as a second step, to set up different data management policy with regards to these three types of data, and to store them differently, so that some of the more critical data, like the transaction data, we actually make sure that they are backed up on a daily basis, whereas the analytics and the diagnostic data, these are data which could be actually be filtered and dropped off to some degree during edge analytics, and only the critical ones are backed up into the central data's centre for further processing. This is the second step that we took.

Then last but not least, we also understood that, again because of the transient nature of this application, we need to make the infrastructure application-ready, meaning that it is easy for the end users to be able to set up a IoT instance on demand to measure a certain transaction and to switch it off or to turn it off when it's not required, without having to understand how to set up a VM, how to install a server, how to install storage, etc. These are tasks that you do not expect your hotel branch office users or IT administrator within a hotel to know how to do. So what we did was to standardize an application lifecycle management platform - the multi-cloud orchestration engine that I was speaking of earlier, on to our cloud and extend it to all the branches and to all the users within the hotel, so that now, after they logged in into our portal, instead of having to set up the ESXi, install the OSs, applications and system packages into

the VM and so forth before they can run the IoT applications, they now essentially see an online catalogue. And this catalogue will only present, based on the active directory and the relevant user profile/role/access rights that they have, the suite of applications that they are entitled to use, what price would they be charged on the internal transfer basis, and so on. And it will present itself in the form of catalogue for them to be able to pick and choose to deploy the applications seamlessly.

So in this way, as a third step, we are now enabling the ease of using IoT across a hotel chain, and that helped us overcome the largest inhibitor to adoption.

Anshul Gupta

Thank you Derrick. I think you have made some -- a couple of good points. Sorry, Ashwin, you want to add something?

Ashwin Jaiswal

Yes. I think the moment we mentioned about cloud, it just hit me, the situation when I was coming from Bombay here, I was -- I left home and was trying to reach the airport. We often have experienced latencies when we talked about latency in Microsoft -- I mean, I don't want to name a particular company, but we must have experienced latency in checking our mails. A simple solution, right? There's no urgency on the mails. You can wait for a minute and you can get a mail. Yes? So when you try and access, refresh mail, it takes a few seconds for you to see that mail.

Now just imagine the scenario when it happens in IoT. We're talking about health-critical devices, where the device has to communicate back to cloud, for example, and then to the network and then back to certain devices which are going to take actions. Now, that whole stuff is a cycle. I was just imagining a situation when I was travelling from home to airport, suddenly the entire traffic was asked to stop on the left-hand side completely and we thought there is a emergency. This is like 1 o'clock in the night, midnight. And the entire traffic was pulled over on the left-hand side and suddenly we saw three ambulances running parallel to us, right? And suddenly, when we checked up with the policeman who was standing, he said that there is a heart which is being flown from Kolkata for a heart transplant in Lilavati hospital in Bombay. Now, for that, they had created a channel on the entire traffic route for bypass, right?

Now see what's happening here. I'm sure that the researchers are taking care of these situations and these are going to get tested, but see, there has to be certain kinds of special channels which have to be allocated based on case-to-case basis. We don't know when the patient is going to suffer at home and we're talking about his device communicating to the doctor, or somebody playing a video game probably is transmitting his signals to somebody, you know one of his friend on the other side of it, but that's a game and here it's a life game. How are you going to address that and how are you going to test it? These are all scenarios; it's going to be a humungous transfer anywhere; and especially when you're talking about cloud, so many hops, so

many network connectivities, so many parallel channels, broadband network, wireless network, how are you going to take care of every single thing. It's going to be -- you know I feel a little scared about the situation as we move [to cloud-based]. I think, yesterday when you spoke, the more you destabilise, the more you become confident of doing few more things.

Anshul Gupta

So thanks Ashwin. Those are definitely good points and we have heard from one side of the business, having heard from the service providers and also from Ixia, because you really deal with both sides, the enterprises as well as the CSPs. So I just want to change a gear a little bit and before -- I mean, I just want to quote two examples here. In 2011 there was a scenario in Australia where there was a utility company which suffered a massive security breach, where the SIM cards were really stolen from the smart meters, and they were used for accessing the Internet. And the overall bill shot up to AUD200,000. And what I remember correctly, I think the media quoted that the responsibility was held on to the utility company, not on the service provider (Telstra).

We had a similar scenario in 2011 in Johannesburg where 400 traffic lights were really broken and SIM cards were really stolen from them. And it's about repairing those street lights and that bill went up to $1.3 million. Not only that, of those 400 SIM cards which were stolen, bill for one of the SIM cards was $4,500. 150 SIM cards were actually used before they could really block the other SIM cards. So this really brings forth the importance of testing. Are enterprises really ready and proactive about testing. So I would like to hear your views, Naveen, on this.

Naveen Bhat

Okay. That's a pretty big change of gears. What we observe on the enterprise side is that there is a fairly deep set of pockets to invest in security. However, in spite of the increasing amount of investment in security, the confidence in the chief security officers, or CIOs is still quite low. We also talked about -- yesterday we talked about this multiple times, that we need to go from a reactive mode to a proactive mode on our security posture. And we still do not see that, except in some of the leading Fortune 500 type companies. Many of the other enterprises are still waiting to see -- you know, for something to happen and they'll take action at that point in time.

The sad part of this whole situation is, in these cases of security breaches, which happen, most likely due to inadequate testing -- the testing could be of different types; network testing, pen testing, whatever -- when these breaches happen, the estimate of the losses is not known for a very, very long period. So if somebody breaks into your house and they steal something, you know, at the end of the day, okay, somebody came in, they stole some gold, some TVs, some money, whatever, and you can estimate the losses. If somebody steals something from a shop you can estimate the losses. But when there is a cybercrime of this sort, where SIM cards are stolen, identities are stolen, passwords are stolen, bank information is stolen, the losses can go on for an extended period of time, because that information, first of all, is used by multiple groups of people. And each group sells some of that information to the next

group for a higher price. And then that information is utilised by others and the consequence of the loss and what they call the overhang of that loss can be pretty long. So I think it becomes even more critical for all of the enterprises to start taking really, really active steps, proactive steps to ensure the security of their network.

Anshul Gupta

And I think that's where the regulation is also going to play an important role, in that.

Naveen Bhat

Absolutely. You bring up a very good point. I think, irrespective of what we tell enterprises, unless regulators, whether it's the government as an actual federal body or a ministerial body etcetera, unless they step in and put in regulations that force enterprises to report on breaches and the kind of data that was stolen etcetera, it will never become reality, because then commerce and commercial realities get in the way and sometimes people say, okay, it's great to do all of this but I can't afford it.

Anshul Gupta

Yes. Thank you so much. So we are running out of time, but I think I would just like to check with the audience if they have any questions for the panel. Okay. No questions? Thank you so much. Thank you Manek.

Manek Dubash, NetEvents

Thank you.

[End]