

Message

From: Chao, Henry (CMS/OIS) [REDACTED]
Sent: 9/26/2013 10:51:10 AM
To: Booth, Jon G. (CMS/OC) [REDACTED]; Linares, George E. (CMS/OIS) [REDACTED]; Akhtar Zaman [REDACTED]; Coutts, Todd (CMS/OIS) [REDACTED]; Outerbridge, Monique (CMS/OIS) [REDACTED]; Patel, Ketan (CMS/OC) [REDACTED]; Fasching, Laura [REDACTED]; Schwarzkopf [REDACTED]; Um, Peter (CMS/CTR) [REDACTED]; Thurston, Robert (CMS/CTR) [REDACTED]; Margush, Doug C. (CMS/OIS) [REDACTED]; Skinner, Dennis R. (CMS/OIS) [REDACTED]; Yousef Shahegh (CMS/OIS) [REDACTED]; Royle, Erick B. (CMS/OIS) [REDACTED]; Venkat Basavaraju [REDACTED]; Carter, Cathy T. (CMS/OIS) [REDACTED]; Marc Richardson [REDACTED]; Outerbridge, Monique (CMS/OIS) [REDACTED]; Grothe, Kirk A. (CMS/OIS) [REDACTED]; O'Kussick, James J. (CMS/OIS) [REDACTED]; Murray, Ruairi S. (CMS/OIS) [REDACTED]; Oh, Mark U. (CMS/OIS) [REDACTED]
CC: Winthrop, Monica (CGI Federal) [REDACTED]; Sharma, Hemant (CGI Federal) [REDACTED]; Karlton Kim [REDACTED]; Michael Finkel [REDACTED]; Campbell, Cheryl (CGI Federal) [REDACTED]; Keith Rubin [REDACTED]; Van, Hung B. (CMS/OIS) [REDACTED]; Stephen Walter [REDACTED]; George Schindler [REDACTED]; Stevenson, Corey B. (CMS/OIS) [REDACTED]; Gray, Edward M. (CMS/OIS) [REDACTED]; Donna Ryan [REDACTED]; Snyder, Michelle (CMS/OA) [REDACTED]; Trenkle, Tony (CMS/OIS) [REDACTED]; Igor Rafalovich [REDACTED]; Jason Hase [REDACTED]
Subject: FW: FFM UI PST results & Concens on System Peformances
Attachments: Summary of PST activities for Monday; PST Summary for Tuesday; PST Summary for Wednesday; Update: PST Summary for Wednesday
Importance: High

Just talked to Jon and it looks like the bar for hc.gov on the learn side is set at 10k concurrent users. I would say that post October 1st the peaks will be at 50k or greater and the Get Insured side for consumers will be at least the same and likely more, and then to compound the demand on the internal FFM architecture to support the call center (web services to the online application) plus the fetching person view by Serco all need to be factored in to a meaningful assessment and test of performance and ability to sustain HA without having a meltdown starting at only 10k on the get insured side.

The coordination for demand on the FFM architecture and Hub also needs to have the network (internal to the cloud, external to the cloud to the Internet, and external to the cloud to CMSNet plus the replication to HP) and all other factored computed to some mathematical expression of total and partial part of the architecture theoretical

capacity. Add other overhead such as security, monitoring, OS level and Application level services such as EFT to the total as well to get a full picture.

George and Akhtar will be working with Jon and Ketan and with all the architects and TMRK/Verizon/URS to get me and answer today (the calculated total and partial capacities of the architecture) and then move quickly to coordinate a set of test to stress the systems to find the breakpoints now.

Monica and CGI Team—I need to know which environment will be the place where we will do this performance testing at the peaks I mentioned—that means the infrastructure folks and the architects have to have the configurations set correctly and triple checked for every element that affects availability and performance.

Separately I need Monique to get this in sync with any necessary execution of temporary and longer term agreements to supply and be able to spawn greater capacity on demand and get that ironed out now—I DO NOT WANT A REPEAT OF WHAT HAPPENED NEAR THE END OF DECEMBER 2005 WHERE MEDICARE.GOV HAD A MELTDOWN (THIS IS TO GET YOUR ATTENTION IF I DIDN'T HAVE IT ALREADY. IF you even detect the slightest need for more compute at TMRK or at HP or both then work on a plan to get it temporarily or permanent if needed.

I expect to see something planned and ready to execute by 6-7pm this evening.

Henry Chao
Deputy CIO & Deputy Director,
Office of Information Services
Centers for Medicare & Medicaid Services
[REDACTED]

From: Chao, Henry (CMS/OIS)
Sent: Thursday, September 26, 2013 9:09 AM
To: 'Campbell, Cheryl (CGI Federal)'; Rich Martin [REDACTED] George Schindler
[REDACTED]
Cc: Outerbridge, Monique (CMS/OIS); Oh, Mark U. (CMS/OIS); Van, Hung B. (CMS/OIS)
Subject: FW: FFM UI PST results & Concrns on System Performances
Importance: High

George
Cheryl
Rich,

I am going to raise this at the 9am so be ready with a response on how we will together focus on getting performance up to at least 10k or greater concurrent users without defects, which will stop us from ever finishing performance testing.

THanks

Henry Chao
Deputy CIO & Deputy Director,
Office of Information Services

Centers for Medicare & Medicaid Services
[REDACTED]

From: Zaman, Akhtar (CMS/OIS)

Sent: Thursday, September 26, 2013 8:50 AM

To: Oh, Mark U. (CMS/OIS); Margush, Doug C. (CMS/OIS); Thurston, Robert (CMS/CTR); Basavaraju, Venkat (CMS/OIS); Walter, Stephen J. (CMS/OIS); Shao, Lijun (CMS/CPI); Van, Hung B. (CMS/OIS)

Cc: Donohoe, Paul X. (CMS/OIS); Grothe, Kirk A. (CMS/OIS); Outerbridge, Monique (CMS/OIS); Chao, Henry (CMS/OIS)

Subject: FFM UI PST results & Concrens on System Performances

Hi All,

We ran several Performance & Stress Test (PST) cycles (starting from 10 Vuser to 2000 Vuser) since last 3 days but the results are not good and not consistent at all (see attached PST results between 09/23-09/25). The only success we had on **Monday (09/23)** when we ran a 500 VUser cycle and received acceptable response time (<5 sec) for all the FFE individual application transactions except **Race & Ethnicity** which took about 65 sec. Pls note that our focus was limited to Individual application submission, and since we didn't get the appropriate eligible results (stuck with "Almost finished" issue), we couldn't continue to Plan results & enrollments. During this test only 1 [REDACTED] was used in AZ.

CGI identified the high RT for the Race & Ethnicity and applied a hotfix – which I believe mainly fine tune the [REDACTED] associated with the Race and Ethnicity step. Then we repeated the same 500 Vusers test on **Tuesday (09/24)**, and received good results – the **Race & Ethnicity** took only 4sec. Also Tuesday's test was involved 10 [REDACTED] in AZ. Then we attempted a 2000 Vuser cycle but all the users failed immediately while accessing the Homepage with "Service Temporary Unavailable" error and the test was stopped. CGI determined that some of the requests were not going to the right [REDACTED] nodes. A [REDACTED] issue was routing traffic to a [REDACTED] server which was down.

On **Wednesday (09/25)** we started with the 500 VUser cycle using 10 [REDACTED] configuration – which was an attempt to replicate the results of the successful 500 Vuser test that was run on Tuesday. But the test finished with a lot of errors. The Individual App, Race and Ethnicity Transactions were taking a long time and eventually the system reached a breakpoint, after which everything started failing. CGI said they were investigating a problem in how the messages are being sent across the cluster under load.

CGI decided that they would like to run a test routing everything to 1 [REDACTED] Server to verify that all transactions are working. We ran a 250 Vuser cycle and received mixed results (187 success & 123 failed transactions, mainly login failures). CGI reported that they reviewed the logs from the test and had changed property files and needed to bounce the servers. We ran a 500 VUser cycle, didn't observe any Login failures but system reached a breakpoint within 6 minutes after which the response time started increasing rapidly and users started failing with HTTP Status-Code=502 (Proxy Error). CGI reported that they had identified an issue in MarkLogic and applied another Hotfix by the end of the day.

Later in the Wednesday evening, we ran two more cycles of testing. We attempted two cycles at 2,000 Vusers. Neither was successful. The results were as follows:

Cycle 1

2000 Vuser test (8:40 pm)

Large number of Race and Ethnicity failures.

CGI suspects a load balancing issue. Made changes and restored the database.

Cycle 2

2000 Vusers (9:31 pm)

Saw very high Login response times. As the vusers ramped up there were a large number of login failures. Test was stopped.

We'll have a technical group meeting today (09/26, 3-4pm) to review the recent Performance & Stress Testing (PST) results along with the performance metrics collected and shared by TRMRK & others. We need to summarize what has been achieved so far and what not and the strategy to continue PST in one of our Implementation environment. Also what's our takeaway from the recent PST findings towards the production configuration for 10/01.

Finally the only positive news is the PST between Hub & TDS (SSA, IRS & EES) went pretty good, and we hit the target SLA (93 TPS with SSA, 40 TPS with IRS & 32 TPS with EES) and the average response time was below 2 seconds (about 1.7 sec).

Pls feel free to share any related questions/concerns on the above summary,

Akhtar

This electronic mail (including any attachments) may contain information that is privileged, confidential, and/or otherwise protected from disclosure to anyone other than its intended recipient(s). Any dissemination or use of this electronic email or its contents (including any attachments) by persons other than the intended recipient(s) is strictly prohibited. If you have received this message in error, please notify the sender by reply email and delete the original message (including any attachments) in its entirety.

Message

From: Zaman, Akhtar (CMS/OIS) [REDACTED]
Sent: 9/27/2013 2:35:50 PM
To: Jim Jones [REDACTED]; Deepak Bhatta [REDACTED]; Natarajan, Venkatesan (CMS/CTR) [REDACTED]
Subject: FW: Terremark RedSeal Topology Map

From: Sharma, Hemant (CGI Federal) [REDACTED]
Sent: Friday, September 27, 2013 2:23 PM
To: Chao, Henry (CMS/OIS)
Cc: Outerbridge, Monique (CMS/OIS); Zaman, Akhtar (CMS/OIS); Um, Peter (CMS/CTR); Oh, Mark U. (CMS/OIS)
Subject: RE: Terremark RedSeal Topology Map

Hi Henry,

Here are my observations/reactions to the email from Dave and building on top of Akhtar's email:

1. The scripts are failing so far due to issues like load balancing, inefficient and defective code, and inefficient queries.
 - a. Let me enumerate the primary issues we have run into so far in our performance tests:
 - i. Incorrect setup of the application zone load balancer which caused only one of the two [REDACTED] servers to be recognized. This was fixed for subsequent performance test runs.
 - ii. Inefficient [REDACTED] for person matching. The [REDACTED] was tuned and baselined.
 - iii. The generation of identifiers within MarkLogic was inefficient. This was fixed and verified as part of the 500 user test.
 - b. At this time we have not identified any inefficient and defective code as part of the performance test.
 - c. We tried to run a 2000 user test but ran into issues with the generation of identifies in Marklogic again. We have since fixed this but it needs to be tested in the next performance test.
2. Adding capacity to address bottlenecks like these will likely be ineffective.
 - a. The defects enumerated above are very typical in any performance test and part of the tuning exercise. However, we were able to capture metrics from the tests that indicate the following:
 - i. There is the opportunity to add more capacity to the [REDACTED] servers to support higher volume of users.
 - ii. Each of the [REDACTED] servers will have a certain capacity to support concurrent requests – based on our tests the floor seems to be 250 users per application zone [REDACTED]. We don't have a value for the ceiling, but our best guess would be 500 or so concurrent users. Based on this, the current configuration would likely support 10K concurrent users and we would benefit from an increase in the number of [REDACTED] servers.
3. While a patch has been developed to fix what we believe is the current bottleneck (the Application ID) we currently have no way to test it and will likely run into the next bottleneck far short of what the equipment and VM's are capable of supporting.
 - a. As Akhtar mentioned, we will be testing in the [REDACTED] environment.
 - b. We should be able to verify the fix for the generation of identifiers.
 - c. We may find other issues during our performance testing – we will find and fix these issues as we encounter them. This is a typical performance test cycle.
4. There are also plans to scale [REDACTED] by adding VMs, but currently the environment has only a single VM.

- a. [REDACTED] currently has two [REDACTED], three application [REDACTED], one [REDACTED], four data [REDACTED], and 3 large [REDACTED] for Marklogic.
 - b. Further, we will be increasing the capacity for [REDACTED] to continue with our performance test efforts.
5. This will make it even more difficult to extrapolate and forecast performance once it is operational.
- a. I believe that performance test results can be extrapolated based on results from the [REDACTED] environment for the production environment.
 - b. We will continue to increase the number of users in performance testing as well as the functionality that is tested.

Please let me know if there are any questions of concerns.

Thanks,
Hemant

Hemant Sharma | Vice President, CGI Federal | [REDACTED]
[REDACTED] www.cgi.com

CONFIDENTIALITY NOTICE: Proprietary/Confidential Information belonging to CGI Group Inc. and its affiliates may be contained in this message. If you are not a recipient indicated or intended in this message (or responsible for delivery of this message to such person), or you think for any reason this message may have been addressed to you in error, you may not use or copy or deliver this message to anyone else. In such case, you should destroy this message, and are asked to notify the sender by reply email.

From: Zaman, Akhtar (CMS/OIS) [REDACTED]
Sent: Friday, September 27, 2013 1:36 PM
To: Chao, Henry (CMS/OIS); Um, Peter (CMS/CTR); Oh, Mark U. (CMS/OIS); Sharma, Hemant (CGI Federal)
Cc: Outerbridge, Monique (CMS/OIS)
Subject: RE: Terremark RedSeal Topology Map

I can add my feedback, and Hemant & Peter can correct me and add more

We're planning to move to [REDACTED] from today to continue our performance testing, troubleshooting, tuning etc. Also I believe CGI will move the fixes for the performance bottlenecks in [REDACTED] today and we can verify it as a part of our PST tonight. We'll continue our PST in [REDACTED] for next few weeks

Based on my knowledge, [REDACTED] already have 3 [REDACTED] in each zone (AZ, DZ, & ML) from FFM side and 8 VMs on DSH/HUB side – which we believe is good starting point to continue testing/troubleshooting. And our goal is to bump up the capacity to at least half of the production (say 8-10VMs per zone) asap to continue the 10k concurrent user load testing and ultimately we're expecting [REDACTED] to be full production sizes to reach our ultimate goal of 50k concurrent user testing. Also we're planning for a single VM testing (1/1/1) to establish some benchmark for the concurrent user capacity that we can use as a baseline for estimating the possible no of VMs needed to support 50k concurrent user load testing.

We're working on the PST data strategy using both TH & TDS options to continue our PST. We'll need your support to extend the SSA & IRS testing window (connecting to their Test/Integration env) till we have the PST data loaded into the TH. I believe SSA is fine connecting to their Integration env as of now but not IRS where we need your help. Pls feel free to let me know any related question/concerns. Thx,

Akhtar

From: Chao, Henry (CMS/OIS)
Sent: Friday, September 27, 2013 1:10 PM
To: Um, Peter (CMS/CTR); Oh, Mark U. (CMS/OIS); Zaman, Akhtar (CMS/OIS); 'Hemant.Sharma [REDACTED]
Cc: Outerbridge, Monique (CMS/OIS)
Subject: Fw: Terremark RedSeal Topology Map
Importance: High

Can one of you please correct Dave's assessment and write the reasons why and get back to me in the next hour?

Henry Chao
Deputy Chief Information Officer and Deputy Director
Office of Information Services
Centers for Medicare & Medicaid Services
[REDACTED]

From: Nelson, David J. (CMS/OEM)
Sent: Friday, September 27, 2013 01:02 PM
To: Chao, Henry (CMS/OIS)
Subject: RE: Terremark RedSeal Topology Map

Henry,

From my perspective. This is not a theoretical capacity planning exercise. I am confident your team of architects have thoughtfully worked through capacity sizing (VMs, CPUs, Memory, Storage, Network, etc.) based on assumptions such as efficient code.

The performance test results in [REDACTED] are based on the 10 [REDACTED]. The scripts are failing so far due to issues like load balancing, inefficient and defective code, and inefficient queries. We have not been successful in moving beyond 500 concurrent users filling applications without income verification. Adding capacity to address bottlenecks like these will likely be ineffective. We must give ourselves the ability to work through these tuning issues and at this point we do not have an operational environment for further performance testing.

While a patch has been developed to fix what we believe is the current bottleneck (the Application ID) we currently have no way to test it and will likely run into the next bottleneck far short of what the equipment and VM's are capable of supporting.

We moved performance testing out of Prod Prime before it became Production and this environment is now connected to the Federal Partner production systems. The plan is to move performance testing to [REDACTED] which is not yet operational. There are also plans to scale [REDACTED] by adding VMs, but currently the environment has only a single VM. This will make it even more difficult to extrapolate and forecast performance once it is operational. Shoring up [REDACTED] and connecting it to the hub and the federal partner test environment needs to be a priority.

I realize everything is a priority but... the facts are that we have not successfully handled more than 500 concurrent users filling out applications in an environment that is similarly in size to Day 1 production. We cannot proactively find or replicate actual production capacity problems without an appropriately sized operational performance testing environment. And, we have not even started looking for tuning issues in the plan select and enroll parts of the application.

From: Chao, Henry (CMS/OIS)
Sent: Friday, September 27, 2013 11:13 AM
To: Um, Peter (CMS/CTR); Outerbridge, Monique (CMS/OIS); Oh, Mark U. (CMS/OIS); Margush, Doug C. (CMS/OIS); Sharma, Hemant (CGI Federal) [REDACTED]; Keith Rubin; Karlton Kim [REDACTED]; Mike Finkel [REDACTED]; Linares, George E. (CMS/OIS); Skinner, Dennis R. (CMS/OIS); Thurston, Robert (CMS/CTR); Zaman, Akhtar (CMS/OIS); Patel, Ketan (CMS/OC); Fasching, Laura [REDACTED]; Nelson, David J. (CMS/OEM); Coutts, Todd (CMS/OIS)
Cc: Berkley, Katrina (CMS/OIS); Rhones, Rhonda D. (CMS/OIS)
Subject: FW: Terremark RedSeal Topology Map
Importance: High

I want to use this as a basis for tracing where we would have performance bottlenecks and Hemant can add his analysis along with Akhtar's extrapolation off performance in a smaller sized environment. By 3pm I want to see an analysis and report describing performance using whatever you guys/gals come up with as a basis for measurement for the network from Akamai in and back out on learn and get insured sides, backside Hub to other agencies performance metrics, max/min for each part including shared parts of the architecture, steady state and burstable levels, and then translated to a performance and capacity framework Michelle and Jim can understand.

Todd will make sure this gets translated and packaged for our use and for Michelle and Jims consumption.

Henry Chao
Deputy CIO & Deputy Director,
Office of Information Services
Centers for Medicare & Medicaid Services
[REDACTED]

From: Schankweiler, Thomas W. (CMS/OIS)
Sent: Friday, September 27, 2013 10:43 AM
To: Chao, Henry (CMS/OIS)
Cc: Um, Peter (CMS/CTR); Margush, Doug C. (CMS/OIS); Oh, Mark U. (CMS/OIS); Thurston, Robert (CMS/CTR); Outerbridge, Monique (CMS/OIS)
Subject: FW: Terremark RedSeal Topology Map

Henry,

Just wanted to provide you with an overdue update you on a task that has been outstanding from security for a while now. Attached is the latest version of the RedSeal network map for Terremark. Folks are working each week to further refine the accuracy of the map. In addition, there have been 6 items identified which Terremark is working to resolve or provide more information on. Each of these are flagging as weaknesses against best practice hardening standards. The CISO's office is advised as they are the ones running the RedSeal program.

Tom

From: Sokoly, Joseph A.(CMS/CTR)
Sent: Friday, September 27, 2013 10:36 AM
To: Schankweiler, Thomas W. (CMS/OIS)
Subject: Terremark RedSeal Topology Map

Tom,

Attached is the most recent Terremark topology map that RedSeal has. Let me know if you have any trouble opening it or have any questions.

Thanks,
Joseph Sokoly

This electronic mail (including any attachments) may contain information that is privileged, confidential, and/or otherwise protected from disclosure to anyone other than its intended recipient(s). Any dissemination or use of this electronic email or its contents (including any attachments) by persons other than the intended recipient(s) is strictly prohibited. If you have received this message in error, please notify the sender by reply email and delete the original message (including any attachments) in its entirety.

Message

From: Zaman, Akhtar (CMS/OIS) [REDACTED]
Sent: 9/30/2013 4:09:09 PM
To: Donohoe, Paul X. (CMS/OIS) [REDACTED]; Deepak Bhatta [REDACTED];
[REDACTED]; Natarajan, Venkatesan (CMS/CTR) [REDACTED]; Jim Jones [REDACTED];
[REDACTED]
Subject: FW: FFM Performance Discussion points

fyi

From: Zaman, Akhtar (CMS/OIS)
Sent: Monday, September 30, 2013 4:03 PM
To: Chao, Henry (CMS/OIS); Outerbridge, Monique (CMS/OIS); Grothe, Kirk A. (CMS/OIS); Oh, Mark U. (CMS/OIS);
Walter, Stephen J. (CMS/OIS); Shao, Lijun (CMS/CPI); Basavaraju, Venkat (CMS/OIS); Van, Hung B. (CMS/OIS)
Subject: FFM Performance Discussion points

Hi Henry,

I'm trying to responding your e-mail yesterday regarding the FFM application performances.

Currently we are seeing performance degradation starting around 1100 - 1200 concurrent users, and most of the pages (except few) are responding within 10 seconds at that load. Few transactions/pages are taking longer such as **Application Summary Save, Family & household Summary Save, Race & Ethnicity etc.** – which should be investigated by the development team, and should be brought to the attention of Monitoring & Helpdesk team. As of today, we're only focusing on Individual application performance testing and going all the way to application submission. The plan is to continue testing to ramp up load to 10k concurrent users. Plan compare has not yet been tested. Currently we cannot get to Plan Compare because the plan data not loaded in [REDACTED] yet. Once available we can provide a page by page metrics for the entire application, registration as well as plan compare. Also Call Center and ESW load scenarios has not been added to the current performance testing because those are not available in current [REDACTED] code.

In addition to the FFM individual application performance testing we're conducting FFM-EIDM integration testing for [REDACTED] accounts creation and step-up to [REDACTED]. EIDM applied a hot-fixes last night (2am) on the "[REDACTED]-identification by-pass" which helping us to continue the [REDACTED] accounts PST. I believe we should focus on the [REDACTED] accounts creation performance testing today considering in first few weeks we may see more accounts creation than individual applications. So far we are able to create 2000 [REDACTED] accounts and step-up to [REDACTED] using 200 concurrent user load, and the average response time was below 4 seconds although the login was about 22 seconds.

The test we are running since last few days give us the limitation of each [REDACTED] servers. We definitely seen some improvements with increased capacity (# of [REDACTED], CPU, Java memory etc.) but still observing issues/bottlenecks with [REDACTED] load balancers/[REDACTED] in every run which needs additional investigation. If we do the plain math (based on the current statistics), the system is able to handle up to 1200 concurrent users with 4 app [REDACTED] – which may require about 30 App [REDACTED] to handle 10,000 concurrent users. And as we're only testing half of the total functionalities (excluding Plan results/compare, enrollments, CCR, ESD etc.) we should even apply a 2.x factor in the estimated calculations – which would make about 60+ App [REDACTED] to handle 10,000 concurrent users. I believe from data zone & ML prospect the overall CPU utilization looks much better, and may not require as high number of [REDACTED] as app zone. Bottom-line we need to focus more on application tuning (code, query optimization etc.) rather just increasing the infrastructures, otherwise by the time we shoot for 50,000 concurrent users we may run out

From performance testing prospect we would need to continue performance testing/troubleshooting in [REDACTED] by bumping concurrent users and also adding more VMs as we progressing. At the same time we would need to get the

most recent code in [REDACTED] and continue scripting for all other missing functionalities (Plan results/compare, enrollments, CCR, ESD etc.). We also need to work with Mathematica to get clean PST data loaded into Test Harness so that we get successful eligibility results and continue to plan results/compare. Our goal is to hit 10,000 concurrent users at first, and then keep shooting for the 50,000 concurrency limits.

Also you have asked to get performance metrics from OC & Akamai, and I have received below stats from Ketan. We'll try to analyze those and factor into our performance testing scenarios.

Below recent statistics based on September 26th which one the busiest day on Healthcare.gov which can provide you some details.

Sept 26th Stats:

Total Page views for the day: 3,857,489

Visits: 628,475

Avg Pages/Visit: 6.14

Avg Page Load Time: 3 sec.

Bandwidth consumption for the day: 1 TB.

Busiest Time of day : 11:00 am to 12:00 pm & 6:00 pm to 7:00 pm EST

Highest Concurrent users at the time of the day: 13,651.

During Peak at 6:45 pm bandwidth traffic: 410.46 mb/sec

During Peak at 6:45 pm request traffic: 1,199.9 req/sec

During Peak at 6:45 pm : 104 page/sec

We have bumped our Akamai bandwidth capacity to 50TB/month. We have purchased security features like WAF & Site shield to protect both Learn & Marketplace from DDoS and any application layer attacks. If we have to increase the bandwidth we can always do that. Learn side has no dependency on marketplace from capacity due to Akamai caching. All users will be first coming to Learn and then go to Marketplace. Learn will control all the incoming traffic to Marketplace.

Pls feel free to let me know any related questions/concerns, Thx,

Akhtar

-----Original Message-----

From: Chao, Henry (CMS/OIS)

Sent: Sunday, September 29, 2013 11:12 AM

To: Zaman, Akhtar (CMS/OIS); Outerbridge, Monique (CMS/OIS); Van, Hung B. (CMS/OIS); Oh, Mark U. (CMS/OIS);

Walter, Stephen J. (CMS/OIS); Shao, Lijun (CMS/CPI); Basavaraju, Venkat (CMS/OIS)

Subject: Re: Discussion points

Importance: High

Akhtar,

I expect you to also have these answers as we continue testing deeper and wider and to help explain in business terms whether if we are:

Determining no degradation of performance for a given segment of testing such as plan compare (expected versus actual tolerance for consumer wait times, for example).

Determining where the degradation begins and the rate at which it begins to falter in delivering a function--this should include Web Services with Call center and Serco.

Making recommendations on adjustments to the architecture at the VM level or some other aspect--doing this in conjunction with the architects.

Working with the architects on determining recommendations for a final configured Prod environment to maximize performance based on what we see in current testing environments.

Working with OC and Akamai determining performance metrics and benchmarks including any dependencies from learn to get insured to we services.

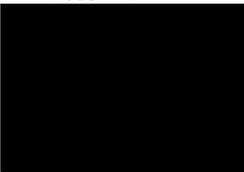
Providing the written analysis and final determinations and estimates for performance in total and by process/function, and where we need to be aware and pre-emptively invoke messaging and/or throttling of users coming in to the online application.

Can you do this starting today (you are doing much of it already) and making some adjustment in how you lead this effort so that by tomorrow midday we have some semblance of a final picture of performance.

Thanks.

Henry Chao

Deputy Chief Information Officer and Deputy Director Office of Information Services Centers for Medicare & Medicaid Services



----- Original Message -----

From: Zaman, Akhtar (CMS/OIS)

Sent: Sunday, September 29, 2013 08:54 AM

To: Chao, Henry (CMS/OIS); Thurston, Robert (CMS/CTR); Outerbridge, Monique (CMS/OIS); 'george.schindler'; 'Cheryl.Campbell';

'Hemant.Sharma'; 'rich.martin';

Van, Hung B. (CMS/OIS)

Subject: RE: Discussion points

I believe Hemant & CGI needs to address those questions based on the Performance testing run last night.

We'll be continue to test today, tomorrow & beyond until we reach at least 10k concurrent users, and eventually 50k. So far, we only covered the Individual application in PST, and today we'll be focusing on scripting Plan results, Plan compare, and enrollments. Also we'll be performance testing FFM-EIDM integration today.

-----Original Message-----

From: Chao, Henry (CMS/OIS)

Sent: Sunday, September 29, 2013 8:43 AM

To: Thurston, Robert (CMS/CTR); Zaman, Akhtar (CMS/OIS); Outerbridge, Monique (CMS/OIS);

'george.schindler'; 'Cheryl.Campbell'; 'Hemant.Sharma';

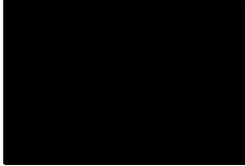
'rich.martin'; Van, Hung B. (CMS/OIS)

Subject: Fw: Discussion points

These are Todd's questions.

Henry Chao

Deputy Chief Information Officer and Deputy Director Office of Information Services Centers for Medicare & Medicaid Services



----- Original Message -----

From: Park, Todd [REDACTED]

Sent: Sunday, September 29, 2013 08:23 AM

To: Chao, Henry (CMS/OIS)

Cc: Snyder, Michelle (CMS/OA)

Subject: Discussion points

Hi Henry, channeling my inner Michelle, a few questions (all of which I think she or I have floated at some point), perhaps for discussion at the 9 am or later today:

-- Has the team run performance/diagnostic testing on the whole FFM, so that we know that the Marklogic bottleneck is in fact the critical, rate-limiting one, as opposed to another bottleneck in, say, Plan Compare or elsewhere, that could also constrain the number of concurrent users?

-- In other words, does the performance testing the team is doing make you confident that the FFM across the board can indeed take 16,000 concurrent users, rising to 60,000-70,000 with the new hardware? - So far we were able to run 2000 concurrent users in [REDACTED], Hemant needs to

-- Are we going to run performance testing today and tomorrow on the growing aggregate collection of hardware (not just one unit of it), so we validate the projections of 16,000/60,000-70,000 with the actual production machinery?

-- Are we testing to make sure that incoming traffic gets properly load balanced across the VMs/units? (This may be accomplished by the previous item)

-- What happens after the 16,000/60,000-70,000 threshold is reached? Is there gradual degradation of response time for users? Rapid degradation? Immediate crashing?

Massive kudos again for the incredible progress the team is making!

This electronic mail (including any attachments) may contain information that is privileged, confidential, and/or otherwise protected from disclosure to anyone other than its intended recipient(s). Any dissemination or use of this electronic email or its contents (including any attachments) by persons other than the intended recipient(s) is strictly prohibited. If you have received this message in error, please notify the sender by reply email and delete the original message (including any attachments) in its entirety.

Message

From: Deepak Bhatta [REDACTED]
Sent: 9/29/2013 7:36:04 PM
To: Akhtar Zaman [REDACTED]; Thurston, Robert (CMS/CTR)
Subject: FW: Venky/Deepak/Jim - Pls review!!!: Discussion points
Importance: High

Hey Thurston,

Can you please help with some of the answers below, especially regarding Akamai.

Akhtar, can you please reach out to OC and get some information as well? Below is what I got so far. Please make corrections as needed.

-----HENRY'S REQUEST-----

Akhtar,

I expect you to also have these answers as we continue testing deeper and wider and to help explain in business terms whether if we are:

Determining no degradation of performance for a given segment of testing such as plan compare (expected versus actual tolerance for consumer wait times, for example).

Currently we are seeing performance degradation starting around 1400 users in the race/ethnicity page. Other pages are responding with SLA (< 10 seconds) at that load. The plan is to ramp up load to 10k concurrent users. Plan compare has not yet been tested. Currently we cannot get to Plan Compare. Once available we can provide a page by page metrics for the entire application, registration as well as plan compare.

Determining where the degradation begins and the rate at which it begins to falter in delivering a function--this should include Web Services-with Call center and Serco.

Call Center and ESW load has not been added to the current performance test. ESD functionality is not currently working. For call center we can add the call center APIs to the current load.

Making recommendations on adjustments to the architecture at the VM level or some other aspect--doing this in conjunction with the architects.

The test we are running today should give us the limitation of each [REDACTED] servers. As the FFM adds more servers today we should be able to determine the capacity of each server.

Working with the architects on determining recommendations for a final configured Prod environment to maximize performance based on what we see in current testing environments.

We will be working with CGI to determine the FFM capability per server and come up with a recommendation. We will be working with the architects from both FFM and HUB to determine the final capacity requirement.

Working with OC and Akamai determining performance metrics and benchmarks including any dependencies from learn to get insured to we services.

Thurston/Akhtar can you help with this?

Providing the written analysis and final determinations and estimates for performance in total and by process/function, and where we need to be aware and pre-emptively invoke messaging and/or throttling of users coming in to the online application.

Our plan is to have 54k applications/hr completed. Please add here as well.

Can you do this starting today (you are doing much of it already) and making some adjustment in how you lead this effort so that by tomorrow midday we have some semblance of a final picture of performance.

Deepak Bhatta | QSSI | www.qssinc.com

ACA – FFM Testing



-----Original Message-----

From: Zaman, Akhtar (CMS/OIS) [REDACTED]
Sent: Sunday, September 29, 2013 12:03 PM
To: Natarajan, Venkatesan (CMS/CTR); Deepak Bhatta; Jim Jones
Subject: Venky/Deepak/Jim - Pls review!!!: Discussion points
Importance: High

Pls see below e-mail from Henry, and I need your help to draft the response by today, pls put high priority on this, feel free to call me for any questions,

-----Original Message-----

From: Chao, Henry (CMS/OIS)
Sent: Sunday, September 29, 2013 11:12 AM
To: Zaman, Akhtar (CMS/OIS); Outerbridge, Monique (CMS/OIS); Van, Hung B. (CMS/OIS); Oh, Mark U. (CMS/OIS); Walter, Stephen J. (CMS/OIS); Shao, Lijun (CMS/CPI); Basavaraju, Venkat (CMS/OIS)
Subject: Re: Discussion points
Importance: High

Akhtar,

I expect you to also have these answers as we continue testing deeper and wider and to help explain in business terms whether if we are:

Determining no degradation of performance for a given segment of testing such as plan compare (expected versus actual tolerance for consumer wait times, for example).

Determining where the degradation begins and the rate at which it begins to falter in delivering a function--this should include Web Services with Call center and Serco.

Making recommendations on adjustments to the architecture at the VM level or some other aspect--doing this in conjunction with the architects.

Working with the architects on determining recommendations for a final configured Prod environment to maximize performance based on what we see in current testing environments.

Working with OC and Akamai determining performance metrics and benchmarks including any dependencies from learn to get insured to we services.

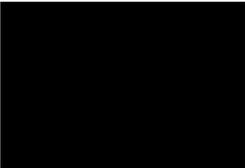
Providing the written analysis and final determinations and estimates for performance in total and by process/function, and where we need to be aware and pre-emptively invoke messaging and/or throttling of users coming in to the online application.

Can you do this starting today (you are doing much of it already) and making some adjustment in how you lead this effort so that by tomorrow midday we have some semblance of a final picture of performance.

Thanks.

Henry Chao

Deputy Chief Information Officer and Deputy Director Office of Information Services Centers for Medicare & Medicaid Services



----- Original Message -----

From: Zaman, Akhtar (CMS/OIS)

Sent: Sunday, September 29, 2013 08:54 AM

To: Chao, Henry (CMS/OIS); Thurston, Robert (CMS/CTR); Outerbridge, Monique (CMS/OIS); 'george.schindler' [REDACTED]

'Cheryl.Campbell' [REDACTED]

'Hemant.Sharma' [REDACTED]

'rich.martin' [REDACTED]

Van, Hung B. (CMS/OIS)

Subject: RE: Discussion points

I believe Hemant & CGI needs to address those questions based on the Performance testing run last night.

We'll be continue to test today, tomorrow & beyond until we reach at least 10k concurrent users, and eventually 50k. So far, we only covered the Individual application in PST, and today we'll be focusing on scripting Plan results, Plan compare, and enrollments. Also we'll be performance testing FFM-EIDM integration today.

-----Original Message-----

From: Chao, Henry (CMS/OIS)

Sent: Sunday, September 29, 2013 8:43 AM

To: Thurston, Robert (CMS/CTR); Zaman, Akhtar (CMS/OIS); Outerbridge, Monique (CMS/OIS);

'george.schindler' [REDACTED]; 'Cheryl.Campbell' [REDACTED]; 'Hemant.Sharma' [REDACTED];

'rich.martin' [REDACTED]; Van, Hung B. (CMS/OIS)

Subject: Fw: Discussion points

These are Todd's questions.

Henry Chao

Deputy Chief Information Officer and Deputy Director Office of Information Services Centers for Medicare & Medicaid Services



[REDACTED]

----- Original Message -----

From: Park, Todd [REDACTED]
Sent: Sunday, September 29, 2013 08:23 AM
To: Chao, Henry (CMS/OIS)
Cc: Snyder, Michelle (CMS/OA)
Subject: Discussion points

Hi Henry, channeling my inner Michelle, a few questions (all of which I think she or I have floated at some point), perhaps for discussion at the 9 am or later today:

- Has the team run performance/diagnostic testing on the whole FFM, so that we know that the Marklogic bottleneck is in fact the critical, rate-limiting one, as opposed to another bottleneck in, say, Plan Compare or elsewhere, that could also constrain the number of concurrent users?
- In other words, does the performance testing the team is doing make you confident that the FFM across the board can indeed take 16,000 concurrent users, rising to 60,000-70,000 with the new hardware? - So far we were able to run 2000 concurrent users in [REDACTED], Hemant needs to
- Are we going to run performance testing today and tomorrow on the growing aggregate collection of hardware (not just one unit of it), so we validate the projections of 16,000/60,000-70,000 with the actual production machinery?
- Are we testing to make sure that incoming traffic gets properly load balanced across the VMs/units? (This may be accomplished by the previous item)
- What happens after the 16,000/60,000-70,000 threshold is reached? Is there gradual degradation of response time for users? Rapid degradation? Immediate crashing?

Massive kudos again for the incredible progress the team is making!

This electronic mail (including any attachments) may contain information that is privileged, confidential, and/or otherwise protected from disclosure to anyone other than its intended recipient(s). Any dissemination or use of this electronic email or its contents (including any attachments) by persons other than the intended recipient(s) is strictly prohibited. If you have received this message in error, please notify the sender by reply email and delete the original message (including any attachments) in its entirety.