

Load Testing Frontier Production

March 15, 2009

Description of the Test

For a base line comparison we used information from Virginia Tech that they provided from their incident in 2007. In the three pages included with this report you will see a large spike in their network and website traffic on April 16. While their website normally averages 129,544 visits a day, that day it was 1,019,966. They average 15GB a day in network traffic (about 3 DVDs worth), but on that day transferred 432GB (about 92 DVDs worth).

At their highest peak Virginia Tech received a little over 300,000 page views in an hour's time, or 83.3 pages a second. After various tests of our load testing software using the backup server for Frontier, we decided on a base setup of one process running on the testing server at its lowest rate. Under testing, that provided us with a page request rate of 142 pages a second (511,200/hour), or 170% that of Virginia Tech's peak load.

We started the test at 6 am on Sunday March 16 and allowed the load to ramp up and achieve a steady state. At that point we allowed the server to run at that rate for 10 to 20 minutes to verify that it could maintain that rate. After that we made a call to Brett Bieber requesting that the normal UNL homepage be replaced by a simple text only emergency page, and then allowed the server to run under those conditions for a while. Brett then changed the UNL homepage back to its original state, and we increased the load on the web server by activating more test servers that requested the UNL homepage.

The Test

Within seconds of starting the test our expected rate had been achieved and no degradation of the service was noted. The load testing software's live statistics reported 0 errors. At that point we raised the load to 700 page views a second. We maintained that rate for 10 minutes. Again, no degradation of service was perceived, and a 0 error status was maintained.

We added the additional testing servers to see the effect they would have on web server performance. Observe the Apache requests per second graph. Note the spike at roughly 6:23 of the rate going from just under 1000 pages/second to almost 2500 pages/second. It was at this point we finally noted that reloading the UNL homepage would take 3-5 seconds longer than normal. We also noted in the testing software's live statistics that the average page loads a second began dropping at a slow rate. Still, no actual errors were reported.

At this point we called Brett and had him switch the page to the plain text version. Immediately the average page views a second began to climb as the lower bandwidth needed to serve out this page allowed more page request per second. Comparing the Apache requests per second, bandwidth usage and server load, you will note that with the change to a text only page the requests per second that were handled rose to 6000 while both the server bandwidth usage and server load dropped. These rates continued until the time we had Brett change the webpage back to the original page.

We declared the test a success at this point and shut off the two additional test servers we had started earlier, leaving just the one initial server running. We would leave it run for a while to build up some long term statistics of the server under a sustained load. The test was finally shutdown at 8:30 am.

Test Results:

Again, using the Virginia Tech statistics as a baseline for our testing, their results were:

Peak Load:	83.3 page views/second
Peak number of hits in a day:	1,019,966 in 24 hours
Peak data transfer in a day:	432GB

Our results:

Peak Load:	700 page views/second (normal home page) (6,000 page views/second -- text only)
Peak number of hits in a day:	88,193,577 in 24 hours (based on 2.5 hr test)
Peak data transfer in a day:	2345GB (based on 2.5 hr test)

On a simple comparison basis the UNL Frontier server is more than capable of handling a load similar to what Virginia Tech faced in 2007.


One concern we did have after finishing the test was that while the server and network were able to handle the increased load, nobody noticed it. The current alert system is designed to only notify someone if a complete failure has occurred. A mechanism should be investigated that would alert someone when a sudden demand is placed on the web server. Another question that arose during discussion of the test results was bandwidth availability at Miller Hall. Currently a major portion of UNL's I1 connectivity is located at SEC. A discussion needs to ensue about whether or not the amount of bandwidth available at Miller Hall is adequate if SEC was unavailable.

Attached Documents:

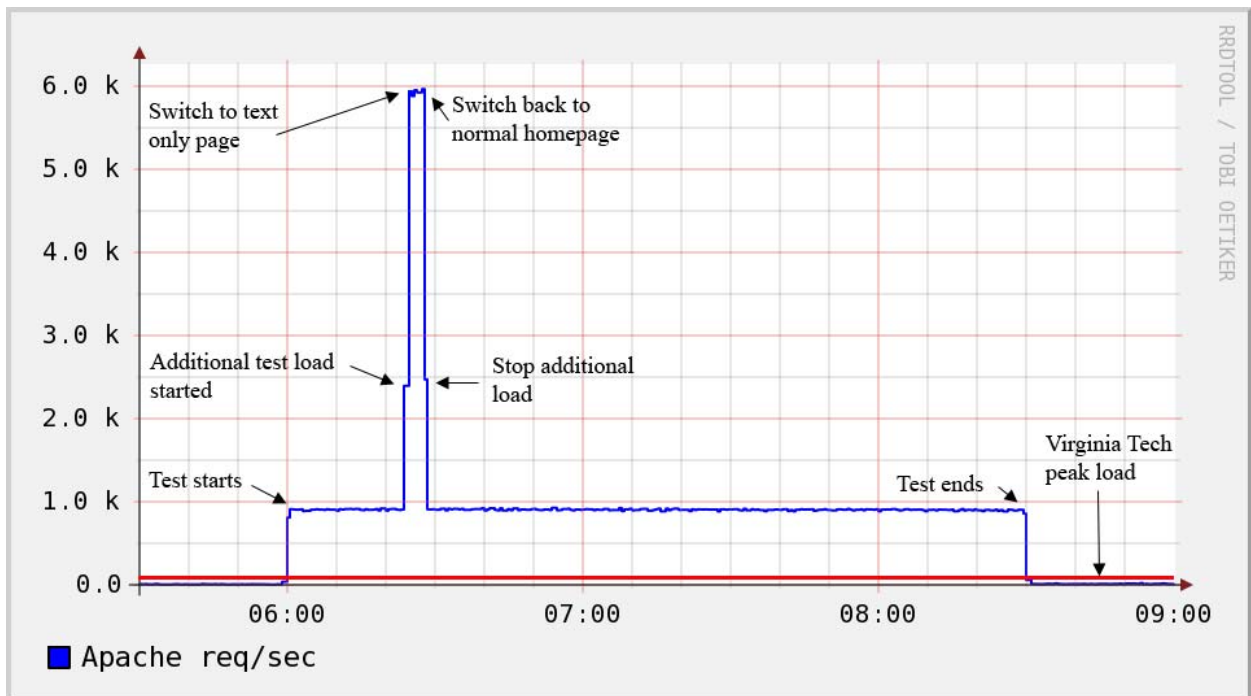
Attached you will find three graphs showing statistics gathered from Frontier showing network bandwidth used, page views a second and server load.

Following that is an analysis of the Frontier Apache server load, showing the total number of hits and data transferred, both for the test period and what it would be for an entire day.

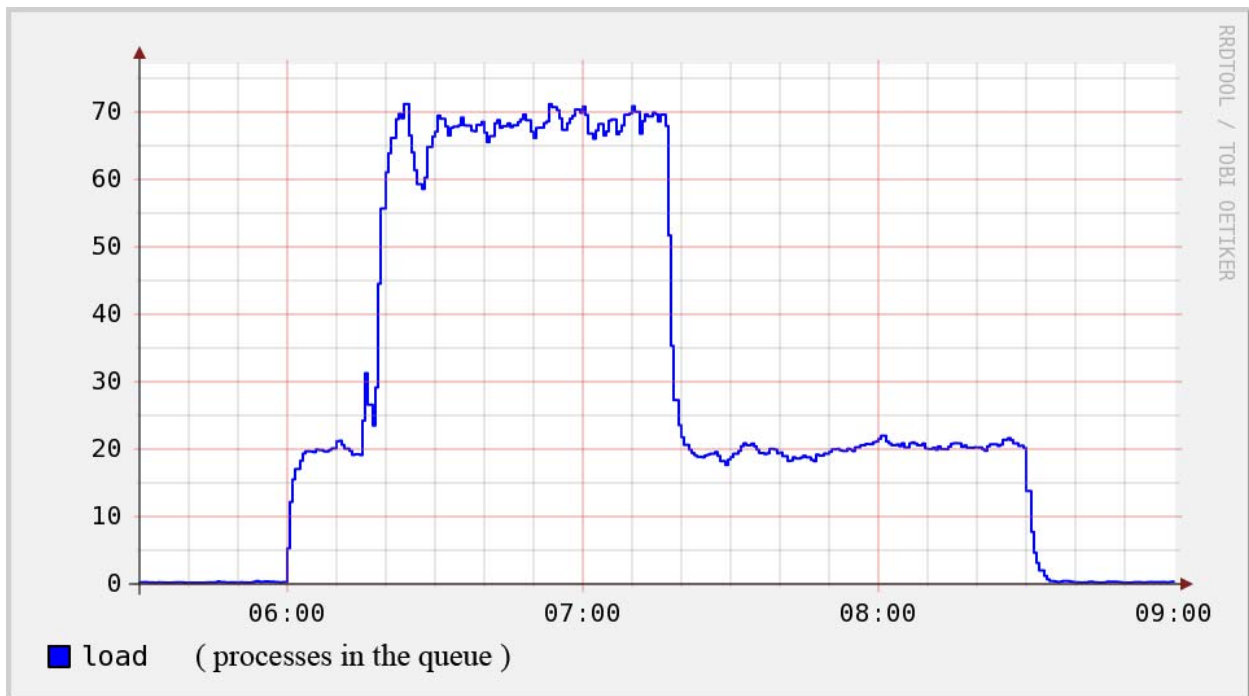
Lastly the three slides from Virginia Tech's presentation at the 2007 HigherEdWeb conference are included highlighting their statistics on April 16 2007.
Frontier Bandwidth Usage:



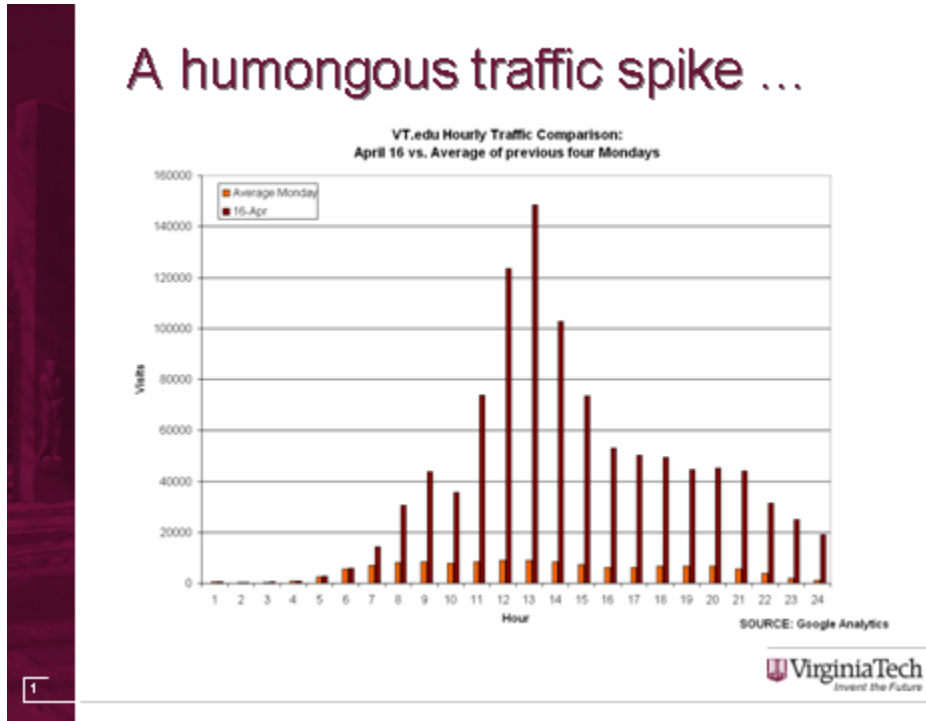
Frontier webpage requests per second:



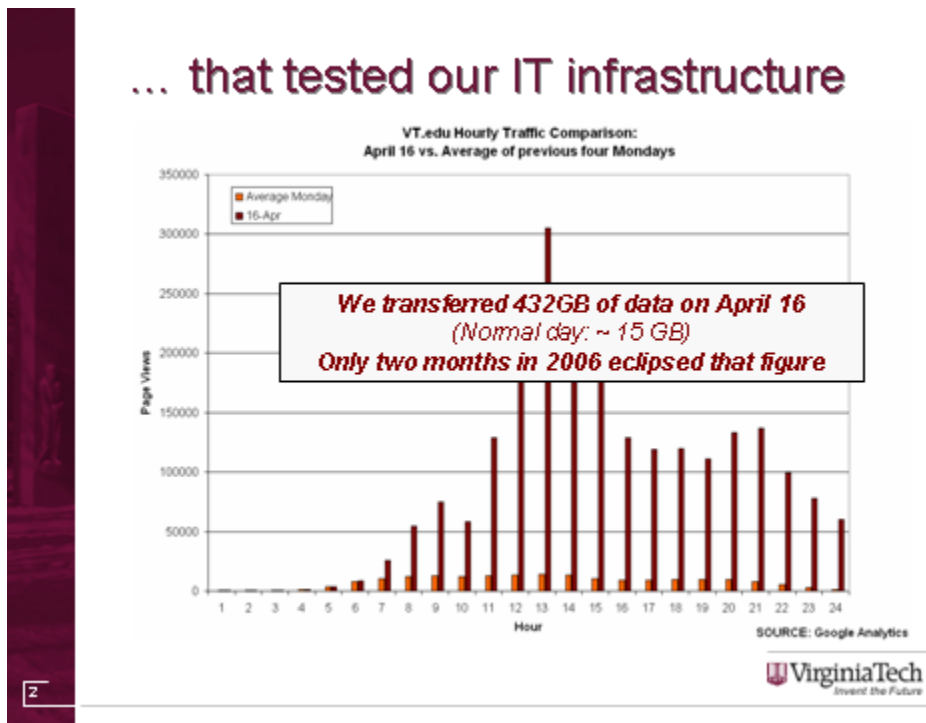
Frontier server load:



Virginia Tech graph comparing average number of visitors to number of visitors on April 16



Virginia Tech graph of average page views an hour compared to the page views an hour on April 16



April 16 traffic summary

Visits

Average	16-Apr	%diff
Monday	1,019,966	687%

Page Views

Average	16-Apr	%diff
Monday	2,294,687	1087%

Trend over time: March-May 2007



Is your IT infrastructure prepared for a spike like this?