

SUPPORTING THE SURGE OF REMOTE WORKERS - 4 STEPS TO ENSURING NETWORK PERFORMANCE



Companies and organizations around the globe are having to shift from an in-office work environment to remote work environment almost overnight. This has caused a big change in the way that people perform their work and engage with co-workers and others. Shifting face to face business, work and personal interactions to online has added significant workloads to digital services and IT systems and infrastructure.

One of the most immediate challenges is making sure that your network performance supports the increase of people using your virtual private network (VPN) and the variety of workloads. How do you determine what you need for doubling, tripling or more the number of users on your VPN?

Are employees holding audio-only meetings? Are they having conference calls that require sharing presentations? Are they doing video conferencing, are they sharing large files?

BMC has the same challenges as other companies around the world. We nearly doubled the number of remote workers in a matter of a few days. This blog shares our best practices using TrueSight Capacity Optimization.

There are 4 basic scenarios that need to be considered to ensure you have the network bandwidth needed to support your remote workers.

1. Modeling your current network bandwidth
2. Analyzing your bandwidth use to detect if expansion is needed

3. Model continuity scenarios – if a network goes down, how can the load be distributed and what does that do to remaining networks
4. Correlate end user response for applications – to determine if slowdown occurs, is it a network issue or compute issue

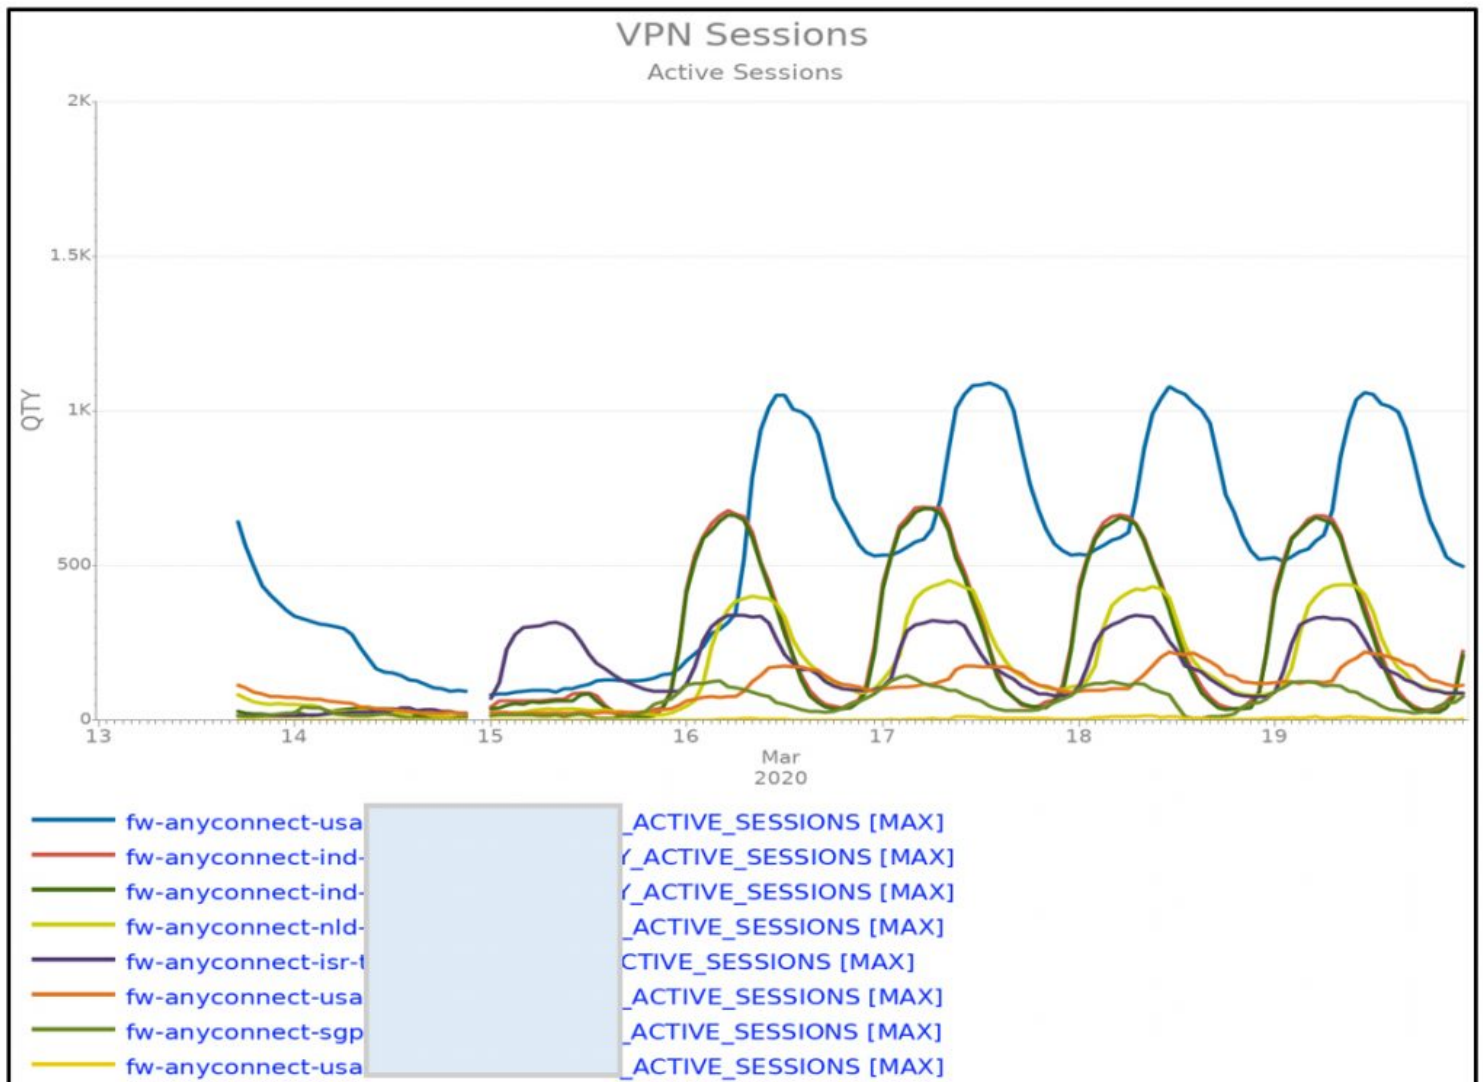
We used the following metrics in performing these 4 scenarios.

- Number of VPN Active Sessions
- Internet Utilization
- Bandwidth of Network Interface
- Link usage per connection
- Input Bit Rate by Network Interface
- Output Bit Rate by Network Interface
- Response Times
- CPU Utilization

1) Model your current bandwidth

The first step to supporting your new remote workers is understanding the current use of your network.

The chart below allows us to gain insight into the number of VPN sessions per location.





You can continue to analyze and visualize active VPN sessions over time to look for trends or potential anomalies.

2) Analyze impact on network and infrastructure

You want to make sure the devices and network servers do not become the bottleneck to good performance.

In the example charts below you can see that there is a significant workload increase on 2 different servers. There is a 100% increase in workload, but server capacity is still only between 20-40%. This is OK for now, but we continue to analyze these devices to understand the capacity increase as you add more users.


Below is a chart showing the normal usage. It shows the increase in usage due as adding additional remote workers and what the new trend in usage will be with this added workload.




Increase in AnyConnect device CPU utilization by 2x, but well within the threshold

Each device is different. It is important to look at the bit rate utilization on those devices. The nature


of different workloads means different consumption rates. For example, an audio only conference will use fewer bits that a video conference. Another example is dependency on AWS cloud services will cause higher network usage. Remote workers may be backing up their laptops to using capacity of network servers along with network bandwidth. Monitoring the capacity of the network ports and network servers to understand resource needs for maintaining expected end user performance.



Increased jump in bit rate at AnyConnect device at some sites, probably due to the nature of the workload, but no cause for concern since these are 1GB Pipes.

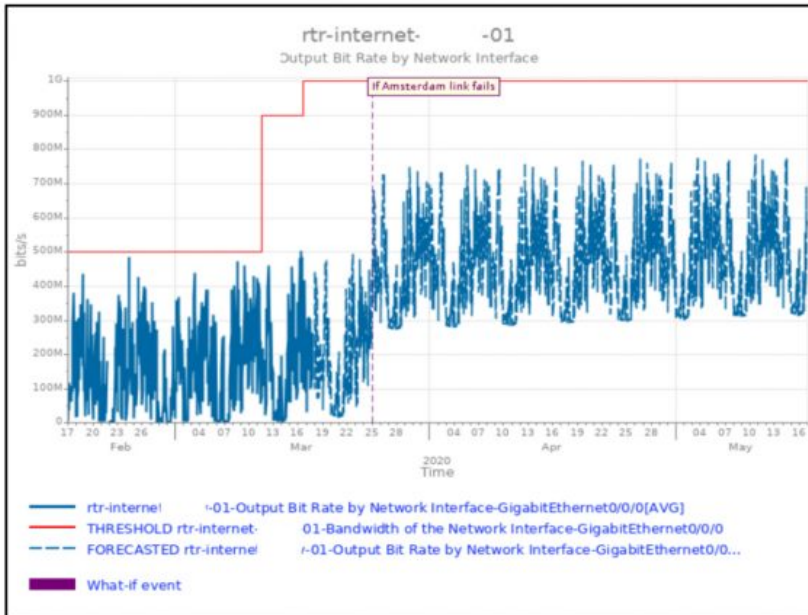


You also want to look at the correlation between active sessions and output bit rate by network interface. This allows you to determine when the circuit will saturate based on the trend for growing number of sessions and current bandwidth.



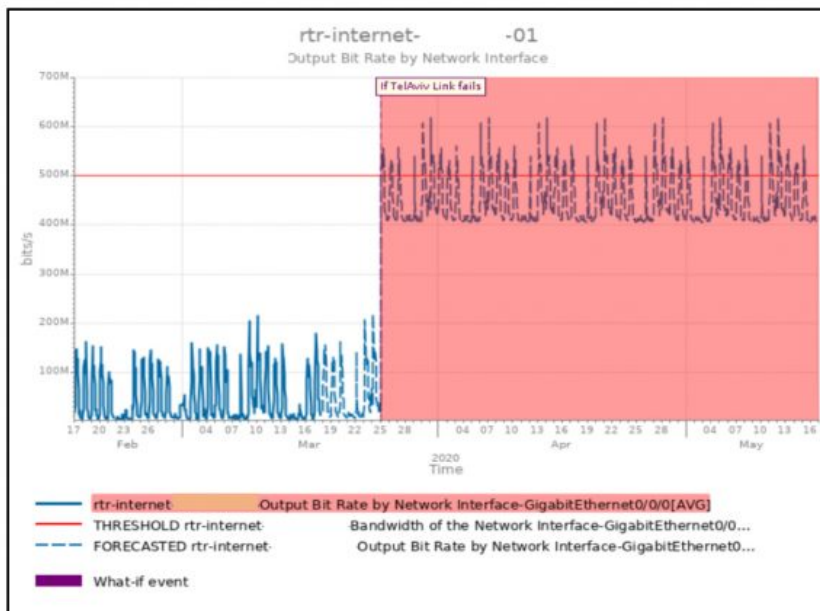
3) Business Continuity

If you have office locations accessing multiple network servers, it is important to understand the workload at the different locations. This is important for understanding capacity requirements for network servers at each location. This is important for understanding bandwidth requirements, but it is also important for developing a continuity plan. The modeling you did in step 1 should provide you with this information. You can perform “what-if” analysis and model continuity scenarios. For example, if the network connectivity at one location failed, how can you distribute the workloads over the remaining networks.



Plan for business continuity. Scenario 1.

If the Amsterdam link fails, the Tel Aviv link can handle the additional load.



Plan for business continuity. Scenario 2.

If the Tel Aviv link fails, the Amsterdam link can't handle the additional load.

4) Monitoring end user response time for an application

Your users have an expectation when it comes to response time. You can model the end-user-response time across the enterprise and ensure that networks are performing as expected.

End user response times for your internal and customer facing applications can be collected from the application performance monitor you are using. This data can be correlated to network response and server response to identify if or when you will have a performance problem.

EBIZ




- EBIZ-Response Time-AUSTRALIA [WAVG]
- EBIZ-Response Time-INDIA [WAVG]
- EBIZ-Response Time-ISRAEL [WAVG]
- EBIZ-Response Time-SINGAPORE [WAVG]
- EBIZ-Response Time-UNITEDSTATES [WAVG]

In the chart below we show that you can add up to 400 users and still maintain current response time. It is important to rerun this model periodically to see if this correlation holds true over time.

ITSM-Response Time vs Concurrent Users

Load VS Load Tot. Volumes - (hourly)



Best practices are to perform regular network performance analysis and modeling on a regular basis, such as weekly. You should develop a “golden model” so you can compare your weekly analysis with the standard that you set for your company.

Want to learn more about network performance analysis. Join the [TrueSight Capacity Optimization Community](#) and ask questions or share your expertise.