Software Defined Networking

How best to optimise the procurement, deployment and management of your network.

By John Bidgood, Systal

Over the last 18 months developments in networking have reflected the trends for separating compute and storage. As such the network has tentatively moved towards software virtualisation of key network and security components using software-defined network (SDN) technology.

Requests for SDN were sporadic at best in 2017 but are now the norm in 2019. The deployment of an SD-WAN Service is the first step to clients utilising an end-to-end Software Defined Environment (SDE). In turn, use of an SDE can lead to the full hyperconvergence of an enterprise's IT infrastructure. Full hyperconvergence is a further reaching project than traditional network virtualisation; it includes the consolidation of network function virtualisation (NFV) with compute and storage at the data centre and remote office/branch into a single set of consolidated and virtualised hardware platforms.

The ultimate goal and best possible outcome is to have a single pane of glass view, allowing you to control all these virtualised platforms from one place.

However, this is a huge exercise in services integration, and one with many steps towards success and numerous associated pitfalls to avoid. Even the first step of deploying SD-WAN can bring considerable benefits but comes with its own risks. For example, when implemented and supported correctly we have seen its potential to reduce client's costs by up to 30% over a traditional WAN service. It has also come with the added benefits of greater security, agility and flexibility to support "multi-cloud" especially when used in combination with traditional internet services.

To maximise these benefits and provide support for long term objectives around using a Software Defined Environment, an innovative approach is needed for the initial procurement, optimisation and support for the Software Defined Network.

In this context, we need to ask two initial questions; the first is whether the traditional model for supporting the WAN with a traditional Internet Service Provider is the

right one? The second question is whether it is better to maintain a hardware virtualisation platform outside of their normal realm of control at the customer premise equipment router?

To answer these questions, we need to first ask how many managed service providers are truly are aligned to your particular goals around the development of a full SDE. Many managed service providers will advise using a traditional ISP, however, many ISPs are aligned to "standard" WAN offerings with support models that are often inflexible to change. This means that they don't take into account the hyperconvergence skills between compute, storage and the need to future proof the whole IT infrastructure, not just the network.

Some managed service providers and ISPs have their own bandwidth solutions, which they will be using as standard to protect their revenue streams. These are often not the most cost-effective solutions for you, so it is important to find a partner who is technology and vendor agnostic, allowing them to assess your current and future requirements against the market and help you find the best solution for your needs.

Another point to consider is geographical coverage. Some ISPs are forced to limit what they can offer and where so that they can cap at the most cost-effective prices. Not all ISPs have assets on the ground they can use; therefore, margin on margin becomes an issue as one ISP buys services from another to then sell on to you. You can save money by working with a partner who can help you procure these services directly.

When things go wrong, will they work pro-actively on anything other than areas within their commercial service control. Think for a moment about the dreaded use case of when the network is reported as slow! Whose responsibility is it to do the root cause

analysis, and will they help if you are finding it hard to pinpoint where the network is slow? This is often compounded by support teams being offshore with no customer network familiarity unless it's a standard ISP service. How many of us run our networks this way?

It's with these challenges, that you should think about the procurement and design of your traditional WAN Services. Therefore, before you extend or source a new WAN contract, we would recommend a methodical approach that does the following:

1. Confirms the end-to-end architecture that the SD-WAN needs to support. Consider where you are hosting applications and how they are being accessed with pertinent security controls. Multi-cloud and remote working has changed and stretched the traditional customer architecture. Additional cloud services locations are like new data centres, therefore complex to migrate and integrate into existing operations and must be approached with care. Security segregation becomes a concern with increasingly remote platforms and the requirements of policies like Bring Your Own Device (BYOD).
2. Confirms what locations need to have 100% network reliability and what can be best efforts. Now, increasing amounts of traffic can be moved online with no tangible drop in performance or reliability. While we agree this represents a trade off on risk considering the cost differences, it remains worth exploring as an option for your business. This is a big trend

enabled by the flexible security mechanisms introduced in SD-WAN.

3. Includes an investment in a CPE platform that supports hyperconvergence of remote site functions, not just the network as part of that architecture. This will reduce TCO on CAPEX and OPEX by helping to hyperconverge the estate into a simplified and standard approach. This will mean that there is no longer a requirement to purchase new hardware each time the organisation wishes to add a new service function to the IT service, it will simply be a case of adding virtual machines with the appropriate licenses. This can be done in minutes rather than wait for the weeks or months it previously took using traditional procurement cycles of purchasing and installing hardware.

4. Decouples the CPE from traditional WAN provisioning models i.e buy unbundled WAN services.

5. Looks across potential multiple ISPs for bandwidth, performance and cloud interconnectivity requirements.

Once you have identified your ideal next steps for your WAN project, the next hurdle to tackle is how to design, deploy and manage the new unbundled services. Given the added benefits of greater control and flexibility, one option is inhouse. Monies saved could be invested on internal staff to run the new virtualised services. Unfortunately, this is not always viable, especially if you only have a few months to go before your WAN renewal, as skills and people take time and money to develop.

This means that enterprises should be on the lookout for potential "integration" partners that can address these challenges on their behalf almost as if they themselves are an extension of the inhouse service model. This integration partner should be technology and vendor agnostic, and able to work on your behalf to secure the best long-term solutions for your business.

An analysis of trends in the industry has shown that you don't have to implement everything on day one to meet your digital workplace and multi-cloud objectives for now and in the future. But by following the right principles early on with the right partnerships, and avoiding vendor lock-in you can, over many years scale a very cost-effective service, rather than undergoing a complete rejuvenation every time it has, inevitably, to change and develop.