WHITE PAPER

Advantages of Managed Spectrum/Wavelength Services Versus DIY with Dark Fiber Model in the Enterprise Market

Introduction

Large enterprises typically buy managed 100G wave service from Communications Service Providers (CSPs) for their business needs, paying monthly OPEX to the CSPs. When the network capacity required by the enterprise is low enough, buying managed 100G wave services is cost-effective for the enterprise. However, as network capacity increases, the monthly OPEX for the managed 100G wave service to the enterprise increases rapidly—compelling the enterprise to seek an alternative business model to contain the cost growth. This leads enterprises to embark on Do-It-Yourself (DIY) models in which the enterprise now procures network system equipment by itself, and only obtains lease or Indefeasible Right of Use (IRU) Dark Fiber (DF) from the CSP to implement the enterprise's transport network. DIY enables the enterprise to control cost. The problem from the perspective of the CSP when their enterprise customers embark on DIY is that the CSP's revenues from the enterprises drop dramatically—since the CSPs are now left only with OPEX from DF or IRU, space, power and colocation as revenues from these enterprise customers. The subject of this white paper is how CSPs can employ managed spectrum service to disincentivize their enterprise customers from embarking on DIY—consequently preventing the CSP's revenues from their enterprise customers from plunging steeply.

For example, Figure 1 shows an enterprise's cumulative cost for managed 100G wave service from a CSP, compared to the enterprise implementing DIY at network capacities of 0.5T, 0.8T, and 1T over a 10-year period.

- At 0.5T network capacity, managed 100G wave service is cost-effective for the enterprise, therefore there is no need for the enterprise to embark on DIY.
- However, as network capacity increases to 0.8T, the cost of managed 100G wave service begins to increase exponentially and intersects the DIY cost curve by year nine, thereby incentivizing the enterprise to embark on DIY with DF.

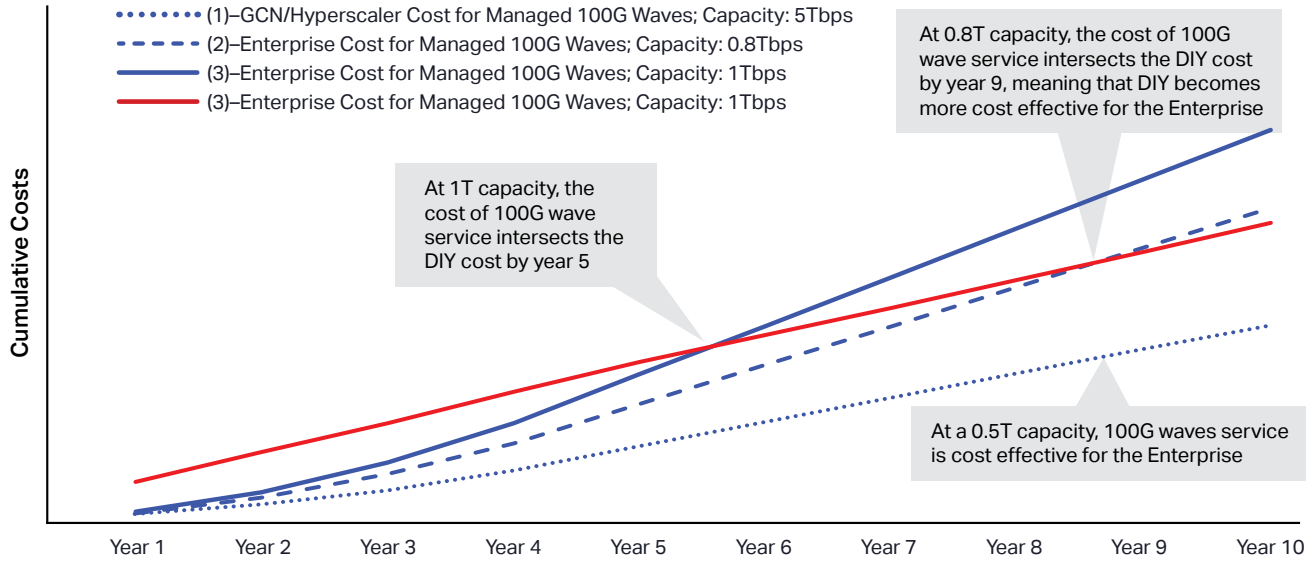

Figure 1. Enterprise cumulative costs for managed 100G wave service versus enterprise DIY for network capacity growth: 0.5T to 1T over 10 years

- As network capacity increases further to 1T, the cost of managed 100G wave service increases more rapidly and intersects the DIY cost curve even sooner by year five, pushing the enterprise to embark on DIY even sooner.

Implementing DIY on DF lease or IRU provides needed network scalability and cost-effectiveness compared to using managed 100G wave service as the enterprise's network capacity grows. However, the greatest challenge for the enterprise is how to operate their DIY networks at the enterprise-grade Quality of Service (QoS) and network reliability of five nines—that is, an availability of 99.999 percent. To achieve a network availability of five nines will require network protection. This means that the enterprise will have to deploy fully meshed or ring networks for protected network operation. In other words, if a large enterprise leases a pair of DF to deploy a network from site A to site B, it will then need to lease a second pair of DF—fully operational and equipped with line systems (Erbium-Doped Fiber Amplifiers [EDFAs], etc.)—on an alternate route between points A and B to divert and reroute traffic in case of network failures.

For large enterprises, implementing DIY with protection is a business imperative

Without the highest level of network resiliency and protection, network failures, blackouts (due to a fiber cut for example), or brownouts (due to network congestion or node failures) pose huge risks for enterprises, with short- and long-term negative impacts on business. These include:

- Immediate lost revenues during network outages
- Loss of peace of mind for enterprise executives
- Poor QoS for enterprise customers
- Reputational damage leading to decreased adoption and unrealized subscriber growth
- Increased customer churn

This white paper shows the challenges for enterprises when the cost of network protection is factored into the economics of implementing DIY. The cost increases significantly—by as much as 42 percent compared to an unprotected network. To meet these challenges, this paper explores managed spectrum service as the most cost-efficient solution, allowing large enterprises to avoid implementing costly DIY with network protection and enabling them to meet CSPs halfway in a risk-sharing partnership. In this arrangement, the enterprise now owns part of the network—the terminals—while leveraging the CSP's vast fiber infrastructure footprint to cost-effectively implement network protection and achieve the highest network resiliency and reliability.

In a similar way, managed spectrum service offers CSPs the opportunity to disincentivize large enterprises from going DIY as the cost growth to them for managed 100G wave service becomes unsustainable. This enables the CSPs to prevent steep revenue declines from their enterprise customers. Managed spectrum service is the solution that offers huge mutual benefits for both the CSP and the enterprise: It maximizes CSP revenues while minimizing enterprise costs—a win-win for both partners.

Overview

The use case: A network has a 40 km link distance with 200G of capacity added per year over five years to reach a total capacity of 1T. The business case presented analyzes four operational models for a CSP to engage with enterprise customers:

Model 1: Managed 100G wave service
A pure OPEX model where the CSP owns and maintains the network from end to end and sells managed 100G wave service to the enterprise customers, who pay a monthly OPEX to the CSP

Model 2: DIY with DF (unprotected)
A model in which the enterprise owns and maintains both terminal equipment and line systems, and leases or IRU DF space and power from CSP

Model 3: DIY with DF (protected)
A model in which the enterprise owns and maintains both terminal equipment and line systems and leases or IRU DF space and power from CSP to implement path protection

Model 4: Managed spectrum/wavelength service (based on Ciena's integrated C&L-band architecture)
A mixed CAPEX/OPEX model in which the CSP owns and maintains fiber infrastructure and line systems, and the enterprise owns and maintains terminal equipment

In this use case, a 100G wave service monthly OPEX of approximately $7,000 USD is assumed. Then, the problems and impacts on CSP revenues and enterprise costs, if an enterprise embarks on DIY, are analyzed. The results show managed spectrum/wavelength service as the ideal business model for both CSPs and enterprise customers, not only because of the aforementioned cost benefits, but it also guarantees the enterprise has a resilient and reliable network and can

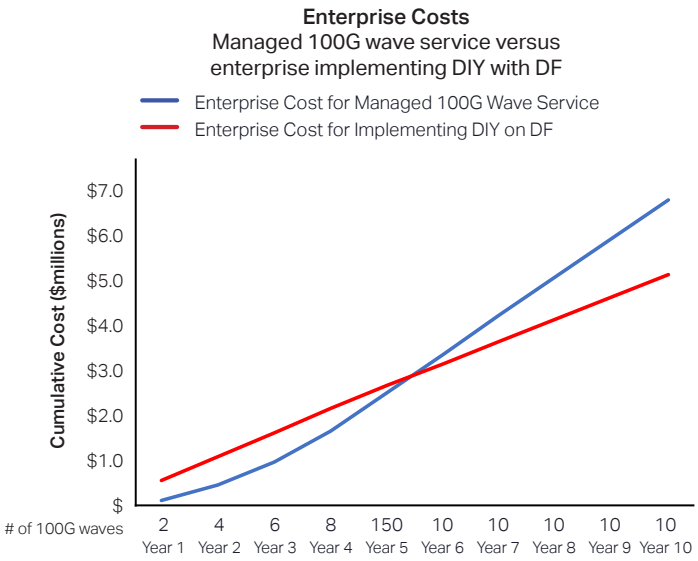

Analysis: At 1T (10 x 100G), cost of managed 100G wave service intersects the DIY cost by year 5—compelling enterprise to go DIY by year 6

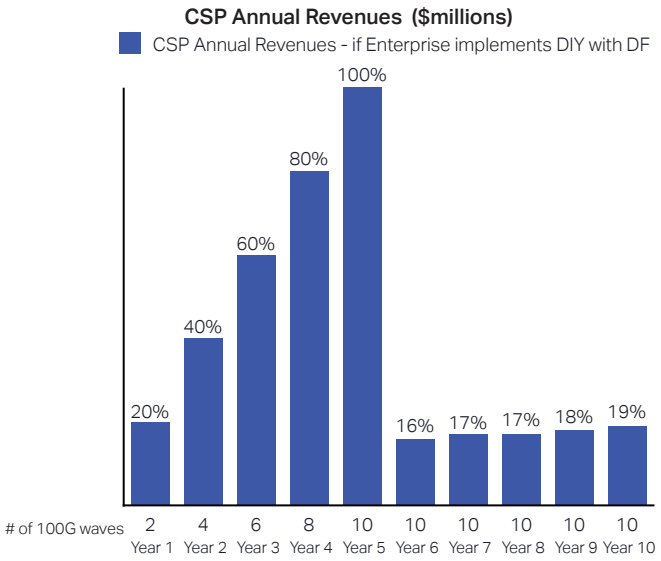

Analysis of CSP revenue (unprotected network): Years 1 to 5 represent the CSP revenues from managed 100G wave service from enterprise. Enterprise embarks on DIY in year 6 and CSP is now left only with OPEX for DF lease/IRU and space/colocation and power from enterprise–leading CSP revenues to plunge by 84 percent between year 5 and year 6

Figure 2. Enterprise cumulative costs and CSP annual revenues from managed 100G wave service at 1T network capacity

cost-effectively achieve an availability of five nines—99.999 percent. A reliability analysis shows that it is impossible for the enterprise to achieve an availability of 99.999 percent by operating a network on a single pair of fiber.

Problem and risk analyses

Risk of CSP revenues plunging if enterprise implements DIY

Figure 2(i) shows the enterprise cumulative cost for managed 100G wave service from a CSP versus enterprise DIY on DF without protection for 1T capacity over 10 years. The x-axis shows the capacity growth over time: 2 x 100G waves are added per, from year one to year five, for a cumulative capacity of 1T by year five, or an equivalent capacity of 10 x 100G waves. The network capacity is constant at 1T after year five, as shown.

At 1T capacity, the cost of managed 100G wave service intersects the enterprise DIY cost by year five, compelling the enterprise to go DIY by year six. The CSP is now left only with OPEX for DF lease or IRU and colocation space and power from the enterprise—plunging CSP revenues by 84 percent between year five and year six as shown in Figure 2(ii). Managed spectrum service is the business model that the CSP can employ to retain its large enterprise customers and prevent its revenues from falling steeply.

Risk of financial losses, negative publicity, and reputational damage to enterprise business if implementing DIY without protection

Figure 3(i) shows enterprise annual costs for managed 100G wave service versus implementing DIY without protection. From year one to year five, the enterprise paid OPEX for managed 100G wave service to the CSP. The enterprise embarks on DIY in year six, incurring a one-time CAPEX for terminals and line systems, ongoing OPEX for terminals and line systems, and OPEX to the CSP for DF lease or IRU and colocation space and power. The impact of the enterprise implementing DIY in year six is to drop the annual OPEX by approximately 44–56 percent by year seven, thereby bending the enterprise cost curve as shown in Figure 3(ii) and leading to a savings of 25 percent over managed 100G wave service by year 10.

However, implementing DIY without appropriate network protection poses serious risks of network reliability issues, since the network cannot be operated at an availability of 99.999 percent. Given the negative business impacts of unprotected networks, deploying a network on a single pair of fiber without protection via route diversity is not a viable option for enterprise businesses. An enterprise would need to deploy DIY with protection, and as shown below, the cost of such a protected network is very high at 1T capacity under consideration.

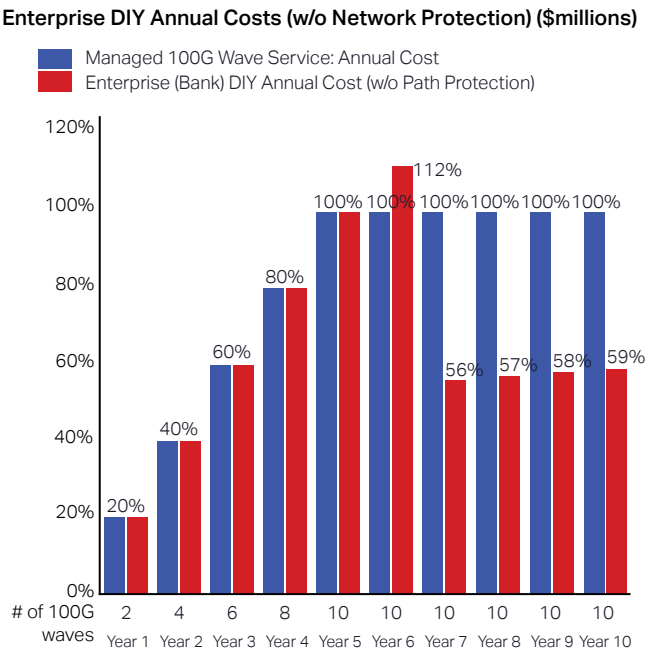

Analysis of enterprise annual costs (unprotected network): From year 1 to 5 enterprise paid OPEX for managed 100G wave service paid to CSP. Enterprise embarks on DIY in year 6, incurs a one-time CAPEX for terminals and line systems and ongoing OPEX for terminals and line systems and OPEX to CSP for DF lease/IRU and space/colocation and power.

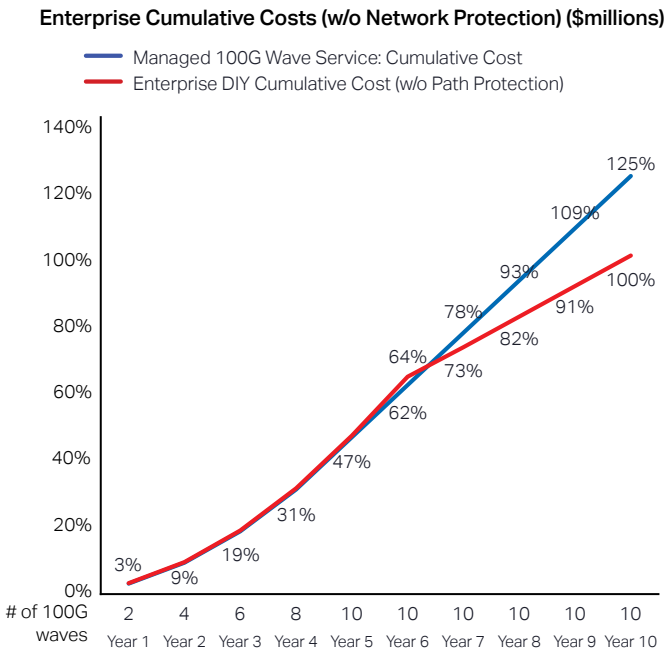

Analysis of enterprise annual costs (unprotected network): Analysis of enterprise implementing DIY in year 6 is to drop the annual OPEX by approximately 44 percent by year 7, thereby bending the enterprise cost curve as shown and leading to a 25 percent savings verses managed 100G waves service by year 10.

Figure 3. Enterprise costs: Managed 100G wave service versus enterprise DIY (without protection) at 1T network capacity

Risk of very high cost to enterprise business if implementing DIY with protection

Figure 4(i) shows the enterprise annual cost for managed 100G wave service versus implementing DIY with protection. From year one to year five, the enterprise paid OPEX for managed 100G wave service to the CSP. The enterprise embarks on DIY in year six, incurring a one-time CAPEX for terminals and two-line systems, ongoing OPEX for terminals and line systems, and OPEX to the CSP for two pairs of DF lease or IRU and colocation space and power. The one-time CAPEX that the enterprise incurs in year six is 84 percent higher than the OPEX for managed 100G wave service, and the ongoing annual OPEX that the enterprise incurs in subsequent years is still higher than the OPEX for managed 100G wave service. Therefore, the impact of the enterprise implementing DIY with protection in year six is to increase the annual OPEX by seven to eight percent over the OPEX for managed 100G wave service by year seven—thereby increasing the enterprise cost curve above that of the managed 100G wave service as shown in Figure 4(ii) and leading to a cost increase of 15 percent over managed 100G wave service by year 10.

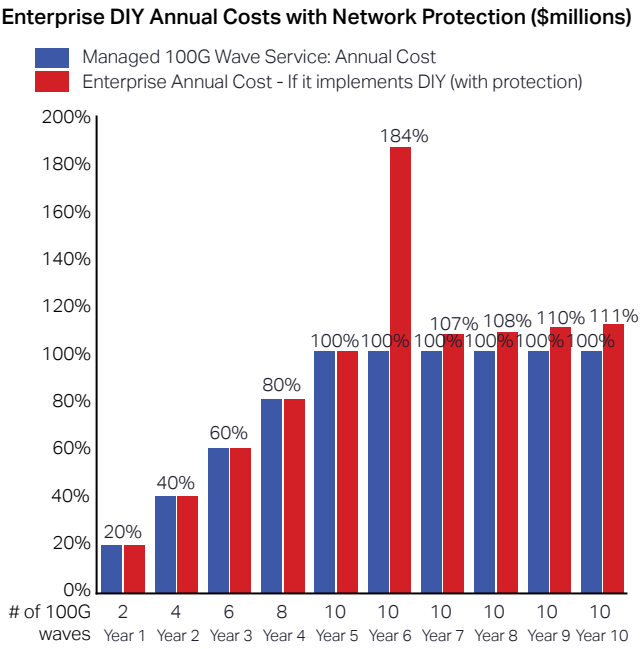

Analysis of enterprise annual costs (protected network): From year 1 to 5 enterprise paid OPEX for managed 100G wave service paid to CSP. Enterprise embarks on DIY in year 6, incurs a one-time CAPEX for terminals and two line systems and ongoing OPEX for terminals and two line systems and OPEX to CSP for two pair of DF lease/IRU and space/colocation and power

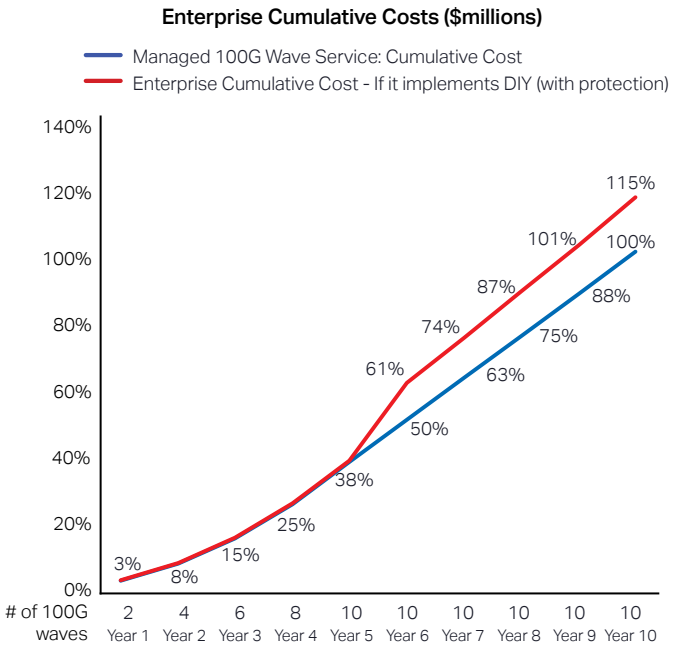

Analysis of enterprise annual costs (protected network): The impact of the enterprise implementing DIY with protection in year six are huge increases in one time CAPEX and annual OPEX that the enterprise cost curve for the protected case is 15 percent higher versus managed 100G waves service by year 10.

Figure 4. Enterprise costs: Managed 100G wave service versus enterprise DIY (with protection) at 1T network capacity

Reliability analysis

$$Availability = \frac{MTTF}{MTTF+MDT} = \frac{MTTF}{MTTF+(MTTD+MTTR)} \quad eqn.1$$

$$Availability = \frac{MTBF-MDT}{MTBF} = \frac{MTBF-(MTTD+MTTR)}{MTBF} \quad eqn.2$$

Where:
MTTF: Mean Time To Failure
MDT: Mean Down Time
MTTD: Mean Time To Diagnose
MTTR: Mean Time To Repair
MTBF: Mean Time Between Failures

As shown by the following examples, without a second fiber pair, it is impractical for an enterprise network to achieve a network availability of 99.999 percent. Assume that in one year (8,760 hours), network maintenance work or failures at a node or terminal bring down the network for just one hour. Using equation two from the reliability analysis, this results in network availability of 99.943 percent, less than 99.999 percent. Again, assume that in a five-year period (43,800 hours), a fiber cut occurs that takes about five hours to repair. This results in network availability of 99.989 percent, still less than 99.999 percent. Hence, the enterprise requires a second pair of DFs on an alternative route—completely operational with line systems—where the enterprise can divert /reroute its traffic in case of network repairs, faults, and failures.

On the other hand, the extensive fiber infrastructure owned by the CSP enables it to redistribute and reroute traffic from a failed network link to other network adjacencies, which typically have enough spare capacity and headroom to take on the capacity from a failed link. This ensures that the traffic reaches its destination via alternate routes while the failed link fault is diagnosed and repaired. It would be a very costly proposition for an enterprise to achieve the same level of network protection as the CSP. It is extremely difficult to envision any scenario where an enterprise can operate the network on a single pair of fiber and achieve a network availability of 99.999 percent. Consequently, the next model shows the business case in which the enterprise adopts managed spectrum service offered by the CSP to cost-effectively achieve network protection with an availability of 99.999 percent.

The solution: A managed spectrum service model to maximize CSP revenue while at the same time minimizing the cost to their enterprise customers

The business case presented in Figures 5(i) to 5(iv) shows four operational models for CSP engagement with their enterprise customers—given a network link distance of 40 km and monthly OPEX for 100G wave service of $7,000 for network capacities of 1T over a 10-year period, based on Ciena's integrated C&L-band architecture and WaveLogic™ 5 Extreme (WL5e) system.

Managed 100G wave service

A pure OPEX model where the CSP owns and maintains the network from end to end and sells 100G wave service to the enterprise, which pays a monthly OPEX

Managed spectrum/wavelength service (based on Ciena's integrated C&L-band architecture)

A mixed CAPEX/OPEX model in which the CSP owns and maintains fiber infrastructure and line systems, and the enterprise owns and maintains terminal equipment.

Managed spectrum—as opposed to shared spectrum—gives the enterprise control of the network at the transponder level for network configurations and maintenance, while the CSP has control of the network line systems, enabling the CSP to implement network protection schemes as needed.

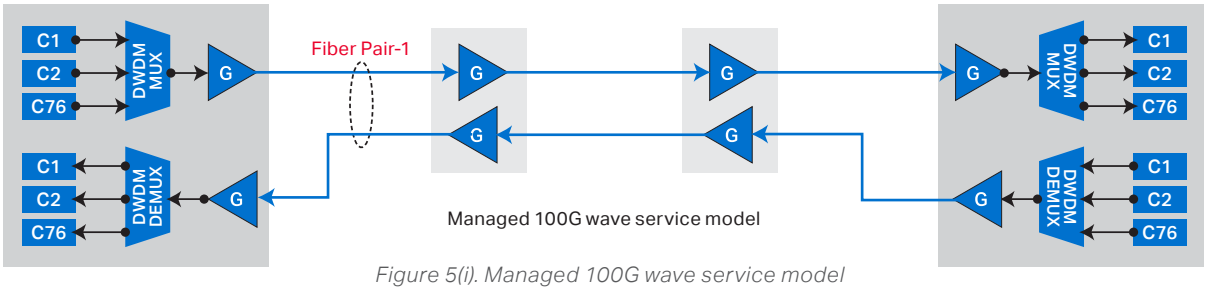

Figure 5(i). Managed 100G wave service model

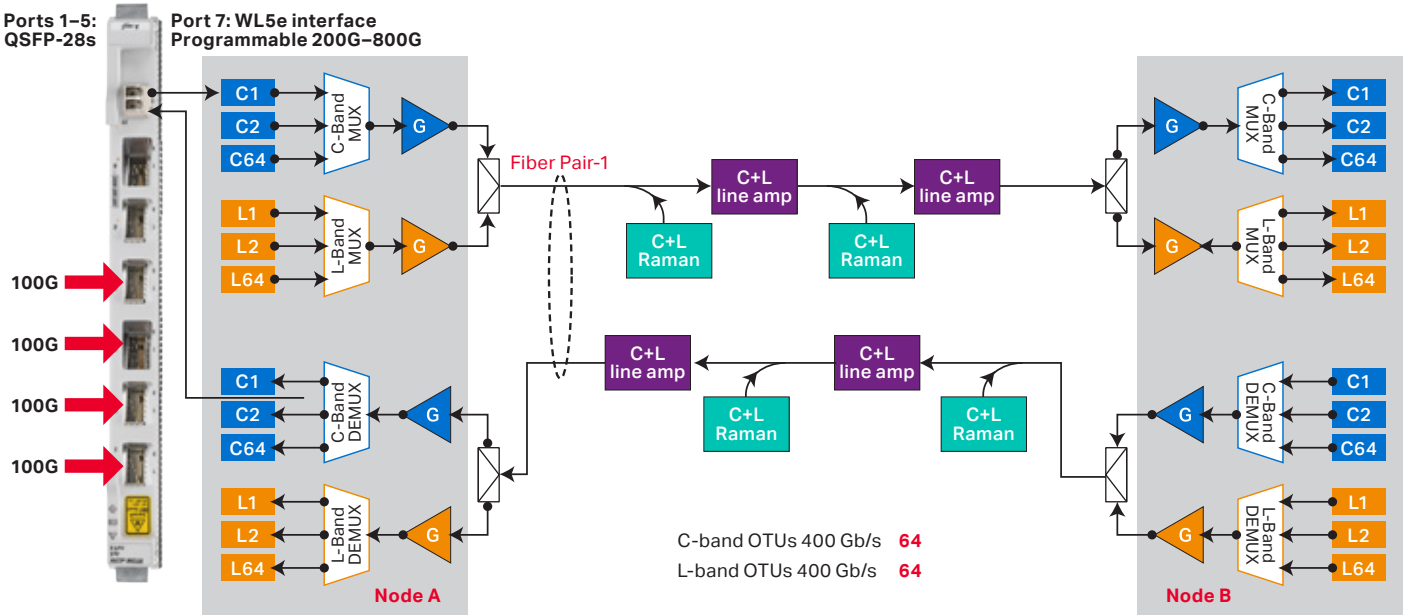

Figure 5(ii). Managed spectrum/wavelength service based on Ciena's integrated C&L-band architecture

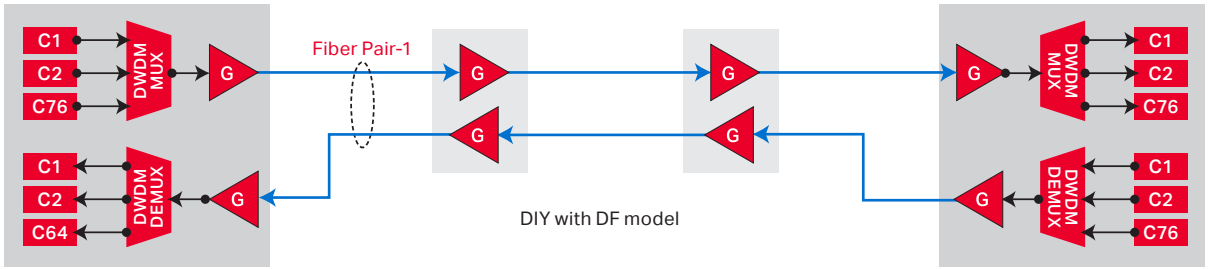

Figure 5(iii). DIY with dark fiber model

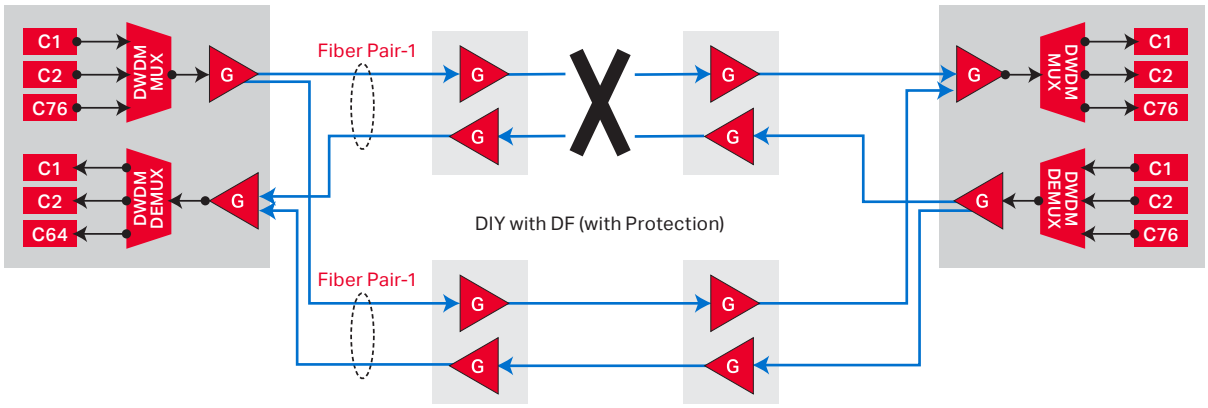

Figure 5(iv). DIY with DF model (with protection)

DIY with DF, without protection

A model in which the enterprise owns and maintains terminal equipment and line systems, and leases or IRU DF, space, and power from the CSP

DIY with DF, with protection

- A model in which the enterprise owns and maintains terminal equipment and line systems, and leases or IRU DF, space, and power from the CSP. For network protection the enterprise needs to lease or IRU two pairs of DF.
- The network with a redundant or ring protection architecture as shown ensures that in the event of a network failure on the main network path—due to a fiber cut or node failure—the signal direction can be automatically reversed to reach the intended destination on the protection path.

Business case modeling results: Impact of managed spectrum service on CSP revenue

Figure 6(i) shows the chart for decision analysis that compares cumulative enterprise costs at a 1T capacity for (1) managed 100G wave service, (2) managed spectrum/wavelength service, and (3) implementing DIY with DF (without protection). As the chart shows, the cost of 100G wave service intersects the DIY cost by year five, meaning that DIY becomes more cost-effective for the enterprise and triggering the enterprise to implement DIY—leading to a steep decline of CSP revenue by more than 80 percent compared to managed 100G wave service by year six as shown in Figure 6(ii). With the managed spectrum service on the other hand, the cost-curve approaches the DIY cost-curve asymptotically but never intersects it—therefore providing the large enterprise no business justification or compelling reason to implement DIY. With the spectrum service, CSP revenues fall by about 31–69 percent compared to managed 100G wave service—50 percent better than if the enterprise were to implement DIY. Given that the managed spectrum service model is a cost-sharing model, the CSP shifts the cost of terminals (CAPEX and OPEX) to the enterprise, for which the CSP receives one-time revenues from the enterprise in year six. This accounts for CSP revenues from managed spectrum service that were 104 percent of managed 100G wave service revenues, as shown.

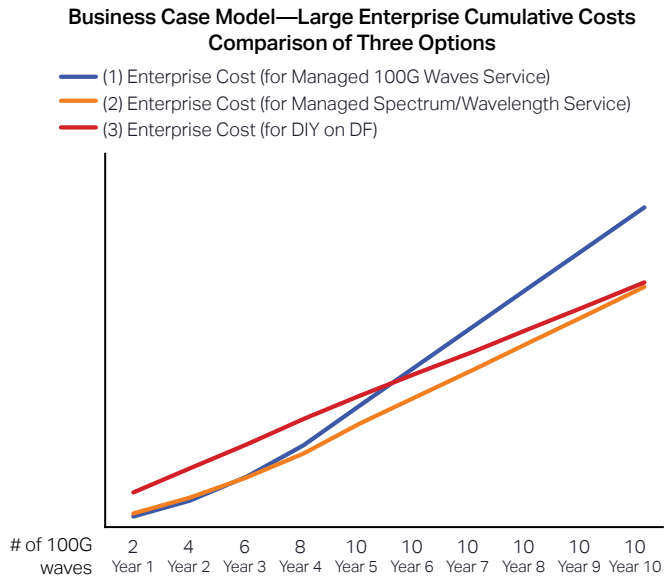

Analysis: Comparison of three options at 1T:
(1) Managed 100G wave service, versus
(2) Managed spectrum service, versus
(3) GCN DIY

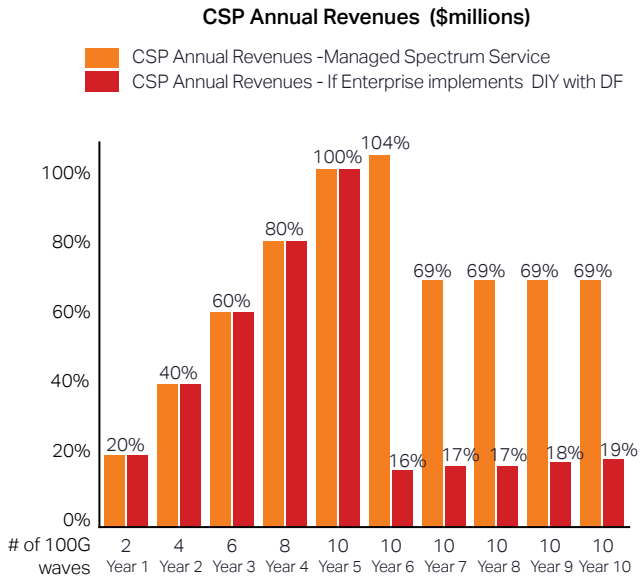

Analysis of impact on annual CSP revenues: A major impact of managed spectrum service is that it prevents CSP revenues from large enterprises from falling off the cliff: Notice that CSP revenue is highest in year 6 (spectrum service introduction) because of the one-time revenue for terminals it received from the enterprise. CSP revenues OPEX for spectrum from enterprise then fall 31 percent to 69 percent in year 7, versus managed 100G waves service. Compare this to the steep decline > 80 percent if enterprise implements DIY.

Figure 6. Comparison of (1) managed 100G wave service, (2) spectrum/wavelength service, and (3) enterprise DIY and impact on CSP revenues at 1T

Analysis of impact on cumulative CSP revenues:

Figure 7 shows the CSP cumulative revenues for the four options normalized to the CSP revenues from the managed spectrum service. By year 10, CSP revenues from the managed spectrum service fall only 18 percent versus revenues from managed 100G wave service. Compare this to the steep decline of CSP revenues of 61 percent if the enterprise implements DIY without protection and a decline of 24 percent if the enterprise implements DIY with protection. Therefore, managed spectrum service lifts and maximizes CSP revenues from large enterprises.

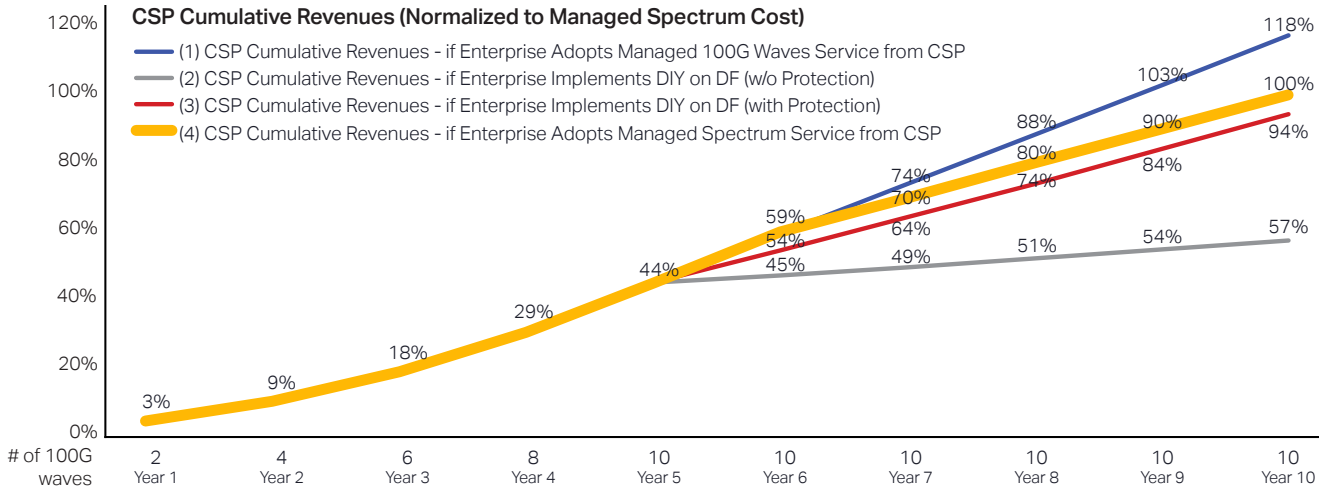

Analysis of impact on cumulative CSP revenues: As the chart shows, the CSP cumulative revenues from managed spectrum service are only 18 percent versus revenues from managed 100G wave service. Compare this to the steep decline of 43 percent if the enterprise implements DIY. The impact is the managed spectrum service will lift/maximize CSP revenues from large enterprises and disincentivize them from embarking on DIY.

Figure 7. Impact of managed spectrum service on CSP revenues at 1T network capacity

Business case modeling results: Impact of managed spectrum service on enterprise costs

Figure 8(i) shows the charts for decision analysis, while Figure 8(ii) shows the enterprise annual cost for DIY with protection and the enterprise cost for managed spectrum service—normalized to the enterprise cost for managed spectrum service. From year one to year five, the enterprise paid OPEX for managed 100G wave service to the CSP. The enterprise embarks on DIY in year six, incurring a one-time CAPEX for terminals and two-line systems, ongoing OPEX for terminals and two-line systems, and OPEX to the CSP for two pairs of DF lease or IRU and colocation space and power.

As shown, the one-time CAPEX that the enterprise incurs in year six for implementing DIY with protection is 78 percent higher than the OPEX for managed spectrum service. Then, the ongoing annual OPEX that the enterprise incurs in subsequent years is still higher than the OPEX for managed spectrum service. Therefore, the impact of the enterprise implementing DIY with protection in year six is an increase in the annual OPEX—even above the OPEX for managed 100G wave service.

Analysis of impact on cumulative enterprise costs:
Figure 9 shows the enterprise cumulative costs for the four options normalized to the enterprise costs for the managed spectrum service. Enterprise costs for implementing DIY with protection is shown to be higher than the cost of the managed 100G wave service by 17 percent, and higher than the managed spectrum service by 35 percent. Therefore, managed spectrum service minimizes the cost for large enterprises.

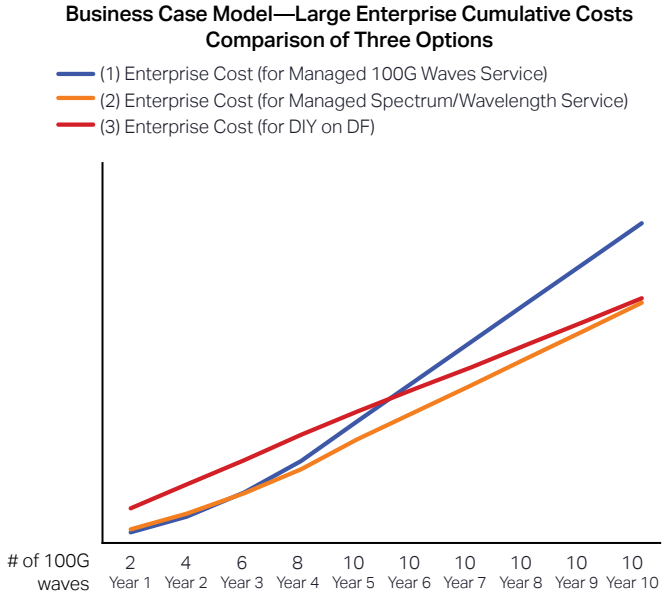

Analysis: Comparison of three options at 1T:
(1) Managed 100G wave service, versus
(2) Managed spectrum service, versus
(3) GCN DIY

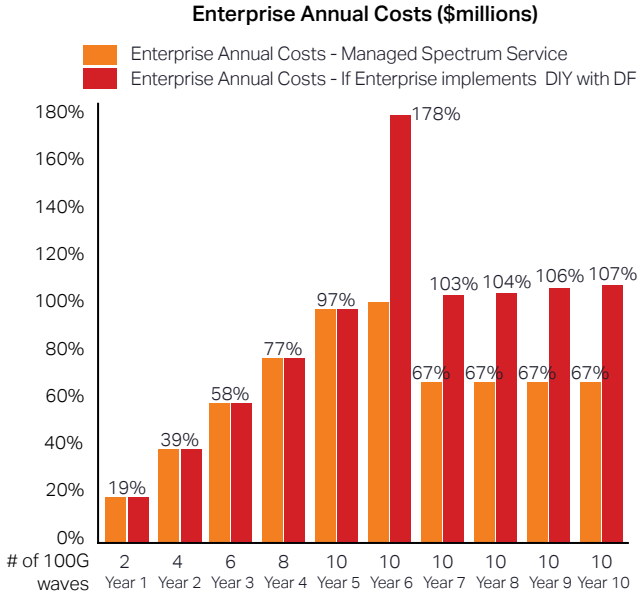

Analysis of enterprise annual costs (protected network): From year 1 to 5, enterprise paid OPEX for managed 100G wave service paid to CSP. Enterprise embarks on DIY in year 6, incurs a one-time CAPEX for terminals and two line systems and ongoing OPEX for terminals and two line systems and OPEX to CSP for two pair of DF lease/IRU and space/colocation and power.

Figure 8. Decision analysis: Comparison of (1) managed 100G wave service, (2) managed spectrum service, and (3) enterprise DIY and impact on enterprise cost at 1T

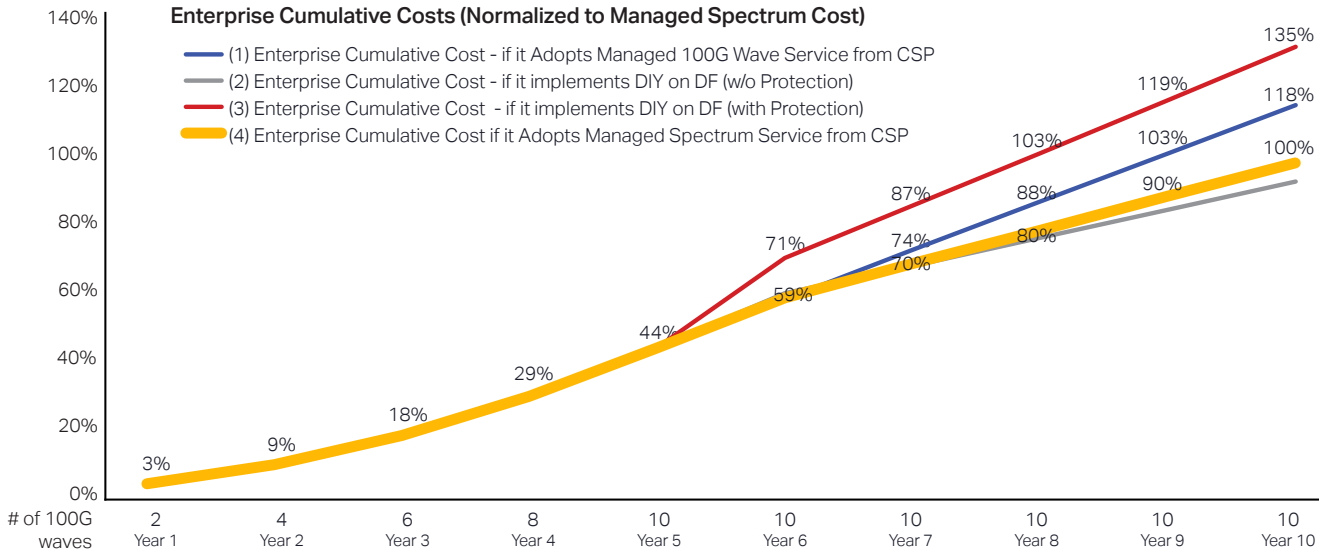

Analysis of impact on enterprise cumulative costs: Managed spectrum service minimizes enterprise costs with a guaranteed nework availability of five nines (99.999%). Enterprise trying to achieve this on their own leads to ~38 percent cost increase.

Figure 9. Impact of managed spectrum service on enterprise costs at 1T network capacity

Conclusions

The four operational models presented in this paper—addressing the use case of a network of 40 km link distance, with 200G of network capacity added per year over five years to reach a total capacity of 1T—demonstrated that managed spectrum service model is the path to success for the CSP to engage with enterprise customers.

Figure 10 shows the impact of each option on CSP revenues and enterprise costs—a 10-year view of the CSP and enterprise cumulative costs. It also shows the risks and opportunities of each operational model:

1. Enterprise buys managed 100G wave service from CSP

- This approach maximizes CSP revenues, but leads to unsustainable cost growth for the enterprise by year five—compelling and incentivizing the enterprise to embark on DIY by year six.
- Risk to the enterprise: High cost as network capacity grows.

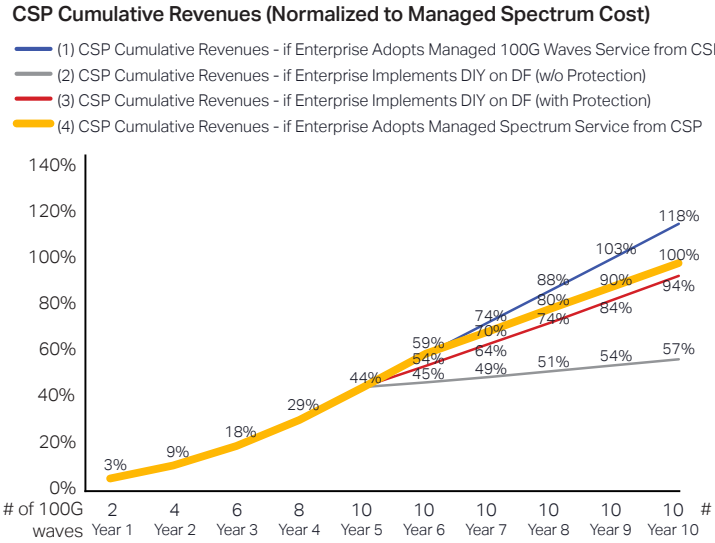

Chart shows that spectrum service maximizes CSP revenues while minimizing enterprise cost with a guaranteed network availability of five nines (99.999%)

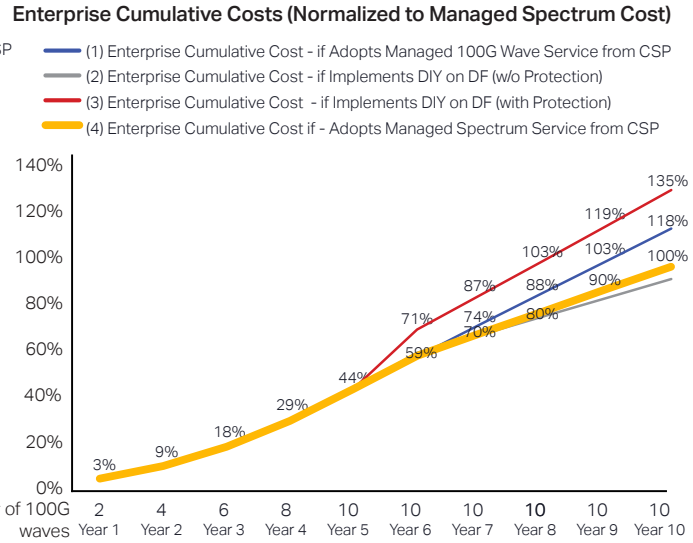

Chart shows that spectrum service minimizes enterprise costs with a guaranteed network availability of five nines (99.999%). Enterprise trying to achive this on their own leads to a ~38 percent cost increase)

Figure 10. Impact of managed spectrum service on CSP revenues and enterprise costs at 1T network capacity

2. Enterprise implements DIY with DF without protection

- Risk to the CSP: CSP revenue collapses, as shown in Figure 10(i). While cost-effective for the enterprise, it lacks the necessary network resiliency and reliability to achieve required network availability of 99.999 percent.
- Risk to the enterprise: Significant negative impact, including:
  - Risk of immediate financial losses during network outages
  - Negative publicity, reputational and brand-name damage
  - Poor QoE and QoS for enterprise customers
  - Lower adoption of enterprise services
  - Unrealized subscriber growth
  - Increased customer churn
  - Loss of peace of mind for enterprise executives

Consequently, this is not a feasible option for enterprises.

3. Enterprise implements DIY with DF with protection

- This approach improves CSP revenues and results in a network with the necessary resiliency and reliability to achieve a network availability of 99.999 percent. However, it is not cost-effective for the enterprise.
- Risk to the enterprise: High cost of implementation and ongoing OPEX—the cost is 17 percent higher than the cost of managed 100G wave service, and 35 percent higher than the cost of managed spectrum service, as can be seen in Figure 10(ii).

4. Enterprise buys managed spectrum service from CSP

- Opportunity for CSP: This is the ideal solution that simultaneously maximizes CSP revenue and minimizes enterprise cost. A managed spectrum service model based on Ciena's integrated C&L-band architecture is the ideal approach that CSPs can employ to disincentivize their enterprise customers from embarking on DIY— enabling CSPs to avoid steep revenue declines. With a spectrum service model, the CSP enters into a cost-sharing partnership with the large enterprises in which the CSP retains ownership of the network infrastructure, including the line systems, but now offloads and shifts the network terminal cost (CAPEX and OPEX) to their enterprise customers. In the cost-sharing arrangement of the managed spectrum service model, the large enterprise now buys managed spectrum service from the CSP as opposed to buying 100G waves.
- Opportunity for the enterprise: This approach implements network protection cost-effectively because resiliency and reliability is inherently built into the CSP network. The extensive fiber infrastructure owned by the CSP enables it to divert and reroute traffic from a failed network link to other network adjacencies, which typically have enough spare capacity and headroom to take on the capacity from a failed link. This ensures that the traffic reaches its destination via alternate routes while the failed link fault is being diagnosed and repaired. It would be very costly for the enterprise to implement the same level of network protection as the CSP.

For large enterprises, implementing DIY with protection is a business imperative—leading to the conclusion that managed spectrum/wavelength service is the solution that results in the most desirable outcomes. This solution simultaneously maximizes CSP revenues and minimizes enterprise costs while achieving the required resiliency and reliability of a network availability of 99.999 percent that the enterprise needs to conduct its business—a win-win for both partners.