FUJITSU

shaping tomorrow with you

Using SFP+ Technology to Manage Network Growth

The Ethernet Dilemma

Bandwidth growth is the most onerous issue in the deployment of Ethernet networks. Typically, a network interface device (NID) is deployed at a new customer site in the form of a ring shared with other customers. At this point, there is a decision to be made: should the NID be put on a 1 GbE ring or a 10 GbE ring? Usually, traffic at the time of deployment warrants only a 1 GbE ring, but based on historical trends, the aggregate bandwidth requirements of this ring will almost certainly increase to warrant a 10 GbE ring in the future. Thus, the service provider has to decide whether to invest in a 10 GbE ring initially without knowing when additional bandwidth will be needed. Alternatively, would it be more appropriate to go with a 1 GbE ring now and change to a 10 GbE ring later? Changing to a 10 GbE ring typically requires changing the NID, an expensive and troublesome activity, but this choice has the advantage of deferring the cost until the bandwidth is needed.

Additionally, customer client ports have grown from 10 Mb to 100 Mb to 1 GbE, and it is easily foreseeable that many will grow to 10 GbE in the near future. Again, what is the best way to meet that client port growth: 1 GbE now or 10 GbE now?

SFP+ to the Rescue

Small Form-Factor Pluggable (SFP) transceivers are popular, small footprint, hot pluggable modules that are available in a variety of capacity and reach options, including 1 GbE. Technology advances now provide an elegant solution to the bandwidth growth issue: 10 GbE performance is available in enhanced Small Form-Factor Pluggable (SFP+) devices that are physically compatible with SFP cages. This new capability provides an exciting and cost-effective solution to common bandwidth growth problems. A NID can be deployed with 1 GbE client ports and 1 GbE network ports using SFPs. Then, when traffic approaches full capacity, 10 GbE SFP+ transceivers can be substituted for the original set. The onerous issue of aggregate bandwidth growth can now be met effectively with a simple change of optical modules and without the burden of changing out the chassis. The same mechanism can accommodate client port growth from 1 GbE to 10 GbE. This solution allows the initial installation to be sized with a more appropriate, lower cost product–1 GbE client and network SFPs–and then grow to 10 GbE when needed. The additional cost is incurred when needed. There is tremendous value associated with this capability, but how much? The rest of this paper will model and quantify the value realized by using SFP+ technology in a typical scenario.

The onerous issue of aggregate bandwidth growth can now be met effectively with a simple change of optical modules and without the burden of changing out the chassis.

The Network Model and Assumptions

This analysis focuses on a pair of core switching devices, each with rings of NIDs at customer sites. The core switching devices are connected with a 10 GbE ring. The network modeled is shown in Figure 1.

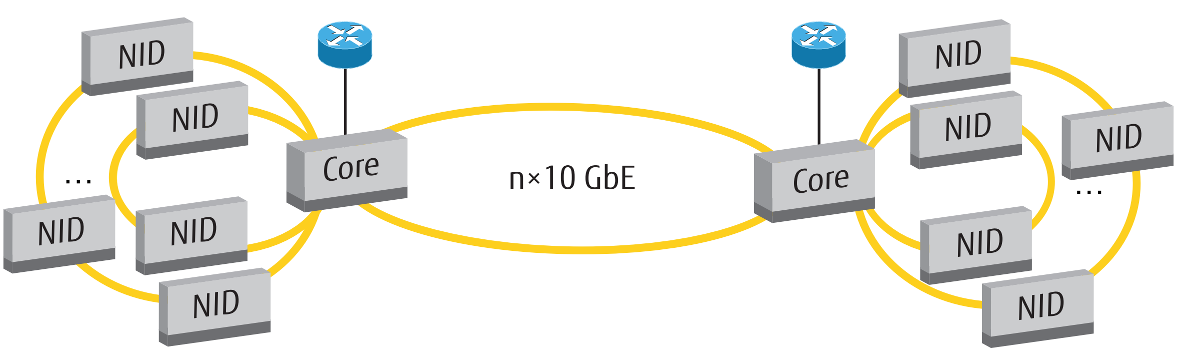

Figure 1: Network model

Using SFP+ Technology to Manage Network Growth

This exercise models the addition of 20 rings to the core nodes each year, with 20 NIDs on each ring. The percentages of 10 GbE rings and 10 GbE client ports added in each year are shown in Table 1. This is not the percentage deployed in each year but a percentage of the total deployment that requires 10 GbE at the end of each year. Meaning, some of the devices deployed as 1 GbE will need to grow to 10 GbE in subsequent years.

	Year 1	Year 2	Year 3	Year 4	Year 5
10 GbE Rings	0%	20%	40%	60%	80%
10 GbE Client Ports	0%	5%	10%	15%	20%

Table 1: Network growth

Scenarios

The network model compares the relative price of access devices when deploying the demand specified for three scenarios.

Scenario 1 minimizes initial cost, Scenario 2 employs SFP+ technology and Scenario 3 minimizes truck rolls.

Scenario 1

Scenario 1 is focused on minimizing initial cost in that the most economical device is deployed at first and upgraded later as needed. Each upgrade requires a truck roll to replace the deployed unit with a higher-capacity substitute. In this scenario, a 2×1 GbE + 2×1 GbE is first deployed. When the network grows to require 10 GbE, the device is replaced with a 2×1 GbE + 2×10 GbE via a truck roll. Then, when the client port needs to grow to 10 GbE, another truck roll will replace the device with a 2×10 GbE + 2×10 GbE.

Scenario 2

Scenario 2 uses a four-port device that offers the flexibility of SFP+ technology. This scenario places the device at customer sites and a switching device at the core, since the latter supports both 1 GbE and 10 GbE. In this scenario, the NID is deployed with 1 GbE SFPs in all four ports. When bandwidth needs to be expanded, an SFP+ is substituted for an SFP, enabling 10 GbE via a truck roll. Although a truck roll is still needed, the SFP change takes much less time than deploying or changing a chassis; thus a truck roll to change an SFP to an SFP+ is less costly than a chassis-swap truck roll.

Scenario 3

Scenario 3, which minimizes truck rolls, involves deploying a device with 10 GbE network ports and 1 GbE/10 GbE–capable client ports at the outset. Truck rolls are minimized, but a more expensive device is deployed initially.

Relative Costs

The relative costs of the devices used in each scenario are shown below.

Scenario	Device	Relative Cost
1	2×1 GbE Client + 2×1 GbE Network	x
1	2×1 GbE Client + 2×10 GbE Network	2x
1	2×10 GbE Client + 2×10 GbE Network	3x
2	2×1 GbE/10 GbE SFP+ Client + 2×1 GbE/10 GbE SFP+ Network	x
2	1 GbE to 10 GbE upgrade	½x
3	2×1 GbE/10 GbE Client + 2×10 GbE Network	3x
2 & 3	SFP Change Truck Roll	y
1, 2 & 3	Truck Roll to Change Chassis and SFP	4y

x = Base hardware price
y = Base truck roll charge

Table 2: Relative costs

This price structure reflects actual product prices, and it allows the model to show the value of using SFP+ transceivers versus reducing truck rolls. The identical cost of the 10 GbE NID in each scenario highlights the impact of truck rolls.

Device Evolution in Each Scenario

In this section, we describe the evolution of the device and its associated relative costs for each scenario as port requirements grow from 1 GbE to 10 GbE.

Scenario 1

In Scenario 1, to minimize costs, the most economical NID—a four-port Gigabit Ethernet device—is deployed initially with 1 GbE SFPs in each port. When the network grows to 10 GbE, the chassis is upgraded via a truck roll, and each network SFP is replaced with an SFP+. Then, when the client ports need to grow to 10 GbE, another truck roll upgrades the chassis and replaces each client SFP with an SFP+. This evolution is shown in Figure 2: the first row shows the evolution of the chassis, and the second and third rows show changes in the client and network transceivers.

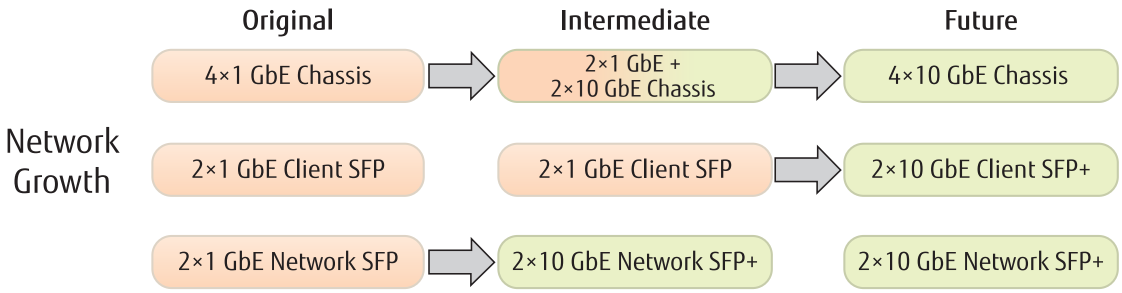

Figure 2: This device growth path minimizes costs.

Scenario 2

In Scenario 2, a four-port 10 GbE device is deployed with 1 GbE SFPs in all four ports. When bandwidth growth requires, an SFP is replaced with an SFP+, enabling 10 GbE. A truck roll to replace the SFP with an SFP+ is less costly than a truck roll to deploy a new device. When the network grows to 10 GbE, the network SFP is changed to an SFP+. Similarly, when the client port needs to grow to 10 GbE, the client SFP is changed to an SFP+. This evolution, shown in Figure 3, deploys the end-state chassis during the initial truck roll and upgrades transceivers as required.

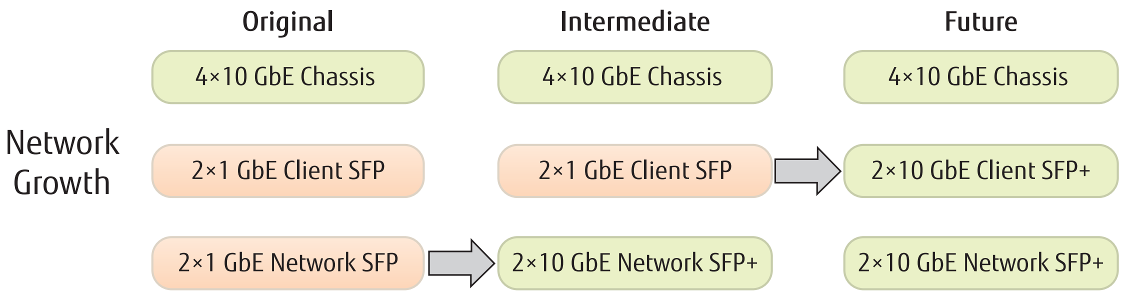

Figure 3: This device growth path employs SFP+ technology.

Scenario 3

In Scenario 3, to minimize truck rolls, a four-port 10GbE device is deployed with 10 GbE network ports and 1 GbE/10 GbE–capable client ports. Thus, there are no changes to the NID or to the ports as bandwidth demands increase, as shown in Figure 4.

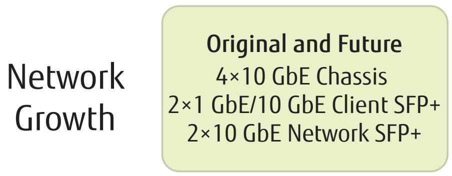

Figure 4: This scenario minimizes truck rolls.

Model Results

We constructed a model of these scenarios to calculate and compare the annual cost under each. The configurations outlined with the devices described in each scenario can be combined with the deployment strategy defined. Then, the enumerated cost of each device and a truck roll is added. After accumulating the total cost of each scenario for each year, the results are shown in Figure 5.

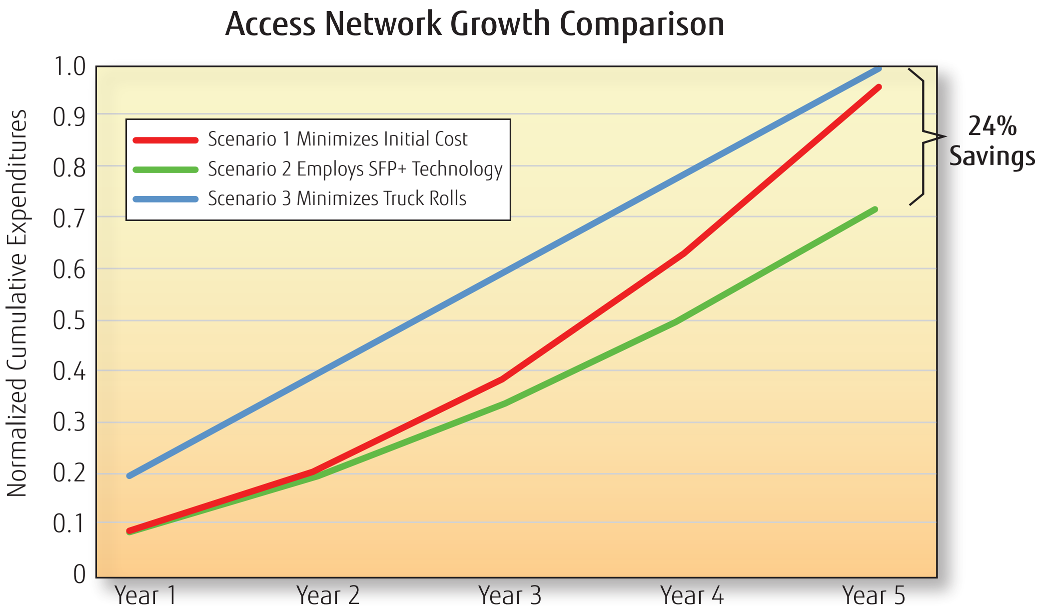

Figure 5: Relative yearly network cost of each device growth path scenario

Summary

The strategy represented by Scenario 3 is not likely to be chosen because it has the highest cost. But the cost-minimization approach represented by Scenario 1 is very common. However, the model described in this paper clearly shows that Scenario 2, which employs SFP+ technology, has greater value than the cost-minimization strategy in Scenario 1. As shown in Figure 5, this study shows a total cost reduction of 24% can be achieved with the use of NIDs that support SFP+ technology. Additionally, capital expenditures are made when needed, not in advance.

Scenario 2 is a flexible, economical approach that employs SFP+ technology. This scenario can be implemented using the Fujitsu FLASHWAVE® 5300 Ethernet Access and Aggregation devices, which are four-port NIDs, and the FLASHWAVE 7120 Micro Packet Optical Networking Platform, a comprehensive core switching device.

This solution provides comprehensive Metro Ethernet Forum (MEF) functionality including eight classes of service (CoS), E-Line, E-LAN, policing, C.8032 sub-50 millisecond protection and Y.1731 performance metrics. Figure 6 depicts this implementation.

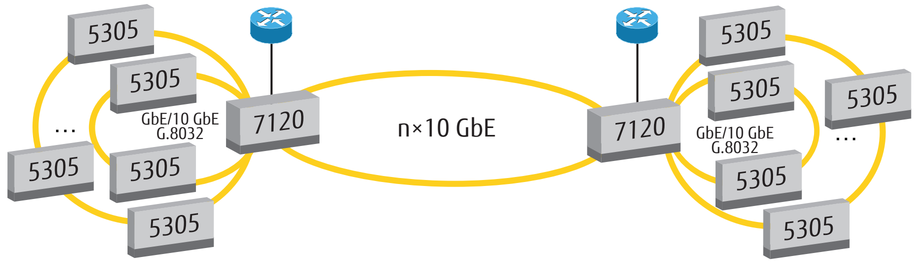

Figure 6: The Fujitsu SFP+ flexible bandwidth access network

Fujitsu Network Communications, Inc.
2801 Telecom Parkway, Richardson, TX 75082
Tel: 888.362.7763
us.fujitsu.com/telecom