

DATA POINT: Entering a period of constant network equipment churn

Many critical network components will need frequent replacements and upgrades

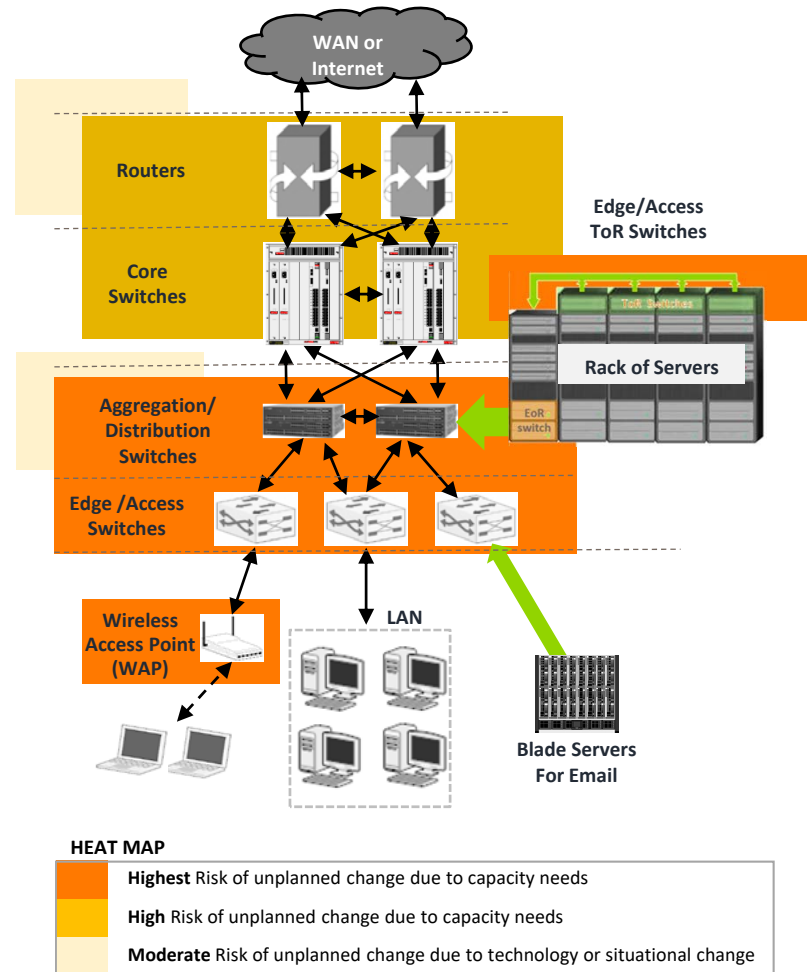
20 FACTORS FORCING CHANGE ON NETWORK EQUIPMENT INVESTMENT STRATEGIES

Of all the areas of technology today, the network infrastructure is under the greatest pressure. Between mobility, data growth, security load, dependence on IT, and pressures on productivity, the network is undergoing forces challenging the traditional longer refresh cycles.

New technologies are forcing expansion and replacement of equipment, and situational factors and uncertainty are creating pressures to hold the assets 'lightly' and to be prepared to abandon current footprints.

Enterprises are just now beginning to see the effects of these, as noted in the 2016 Network Barometer Report by Data Dimension:

“For the first time in five years, networks are getting younger... Companies are starting to refresh equipment earlier in its lifecycle....companies are no longer tactically replacing obsolete devices with like-for-likes, but they’re replacing fully functional and supported aging devices.”



HUNTINGTON TECHNOLOGY FINANCE | Phone: 248.253.9000 | Email: technology-finance@huntington.com

Disclaimer: The information contained in this document is not intended to be relied upon as legal, financial, accounting, tax or other professional advice. Before acting on any information, you should consult with qualified professional advisors and consider your particular objectives, financial situation and needs.

20 FACTORS FORCING NEW APPROACHES	#	IMPACT ON:				
		Core Switches	Core Routers	Edge Switches (ToR, Closet)	Aggregation Switches (EoR, Closet, Core)	WI-FI
100% increase in 40GbE Port counts per year	A	WFR			WFR	
20% Server growth per year	B	WFX		WFR	WFR	
24% growth in traffic within data centers (2014-2019)	C	WFX		WFX	WFR	
25% growth in traffic from data centers to users (2014-2019)	D		WFX	WFX	WFR	WFR
30% growth in 10GbE ports per year	E	WFX		WFR	WFR	
30% Virtualization density increase per year	F			WFR	WFR	
31% growth in traffic between data centers (2014-2019)	G	WFX	WFX			
45% of organizations cite bandwidth requirements as a barrier to cloud use	H	WFX	WFR			
60% of servers in data centers still using 1GbE ports	I	WFX		WFR	WFR	
60% of US network devices are aged or obsolete	J	WFR	WFR	WFR	WFR	
70% of enterprise cloud usage is for file storage	K	WFX	WFX			
74% of WI-FI infrastructure cannot handle mobility	L					WFR
75% of WI-FI access switches cannot support 10GB uplinks	M					WFR
Big Data requires 20GB LANs	N			WFR	WFR	
Move to new Ethernet speeds (25/50GB, new 100GB)	O	WFR		WFR	WFR	
Availability of 40GbE options in converged devices	P	UCN			UCN	
Data center segmentation (DR, growth into colo's)	Q	UCN	UCN		UCN	
Inability to forecast demand growth	R	UCN			UCN	UCN
Availability of SDN (software defined networks) options	S		UCN		UCN	
Availability of NVMe on Fabrics (e.g. FC-NVMe)	T			UCN		

	Will force replacement
	Will force expansion
	Uncertainty in choices and need for nimbleness

#	FACTORS FORCING NEW APPROACHES	DATA SOURCE
A	100% increase in 40GbE Port counts per year	IDC Quarterly Ethernet Switch and Router Trackers, for 2016 Q1 and Q2.
B	20% server growth per year	Infonetics Research, Data Center Strategies North American Enterprise Survey, May 2015, cited in "25Gb Ethernet: The new standard for data center connectivity", DLL/QLOGIC whitepaper 2016.
C	24% growth in traffic within data centers (2014-2019)	Cisco 2015 Global Cloud Index.
D	25% growth in traffic from data centers to users (2014-2019)	Cisco 2015 Global Cloud Index.
E	30% growth in 10GbE ports per year	IDC Quarterly Ethernet Switch and Router Trackers, for 2016 Q1 and Q2.
F	30% virtualization density increase per year	Infonetics Research, Data Center Strategies North American Enterprise Survey, May 2015, cited in "25Gb Ethernet: The new standard for data center connectivity", DLL/QLOGIC whitepaper 2016.
G	31% growth in traffic between data centers (2014-2019)	Cisco 2015 Global Cloud Index.
H	45% of organizations cite bandwidth requirements as a barrier to cloud use	Integra whitepaper Upward Mobility--Moving Business to the Cloud, cited in "Know Your Bandwidth Needs", Electric Lightwave, April 2014.
I	60% of servers in data centers still using 1GbE ports	Rochan Sankar, Broadcom, cited in "The Long Cycles of Enterprise Networks", Timothy Prickett Morgan, Feb 2016, The Next Platform.
J	60% of US network devices are aged or obsolete	Network Barometer Report 2015, Dimension Data.
K	70% of enterprise cloud usage is for file storage	2016 Enterprise Cloud Computing Survey, Clutch.
L	74% of WI-FI infrastructure cannot handle mobility	Network Barometer Report 2015, Dimension Data.
M	75% of WI-FI access switches cannot support 10GB uplinks	Network Barometer Report 2015, Dimension Data.
N	Big Data requires 20GB LANs	"Tracking the Ever-Shifting Big Data Bottleneck", Alex Woodie, Datanami, August 2016.
O	Move to new Ethernet speeds (25/50GB, new 100GB)	Dell'Oro Group Server Report, March 2016, cited in "25Gb Ethernet: The new standard for data center connectivity", DLL/QLOGIC whitepaper 2016.
P	Availability of 40GbE options in converged devices	Product refreshes from Cisco, discussed in "The Long Cycles of Enterprise Networks", Timothy Prickett Morgan, Feb 2016, The Next Platform.
Q	Data center segmentation (DR, growth into colo's)	"Colo's and MSPs cash in on cloud--while they can", David Linthicum, May 2016, InfoWorld.
R	Inability to forecast demand growth	Example: Brocade Network Subscription offering.
S	Availability of SDN (software-defined networks) options	"Top 10 SDN Market Leaders in the Data Center and Enterprise in 2016", Mark Haranas, Feb 2016, CRN.
T	Availability of NVMe on Fabrics (e.g. FC-NVMe)	e.g., Fibre Channel Solutions Guide 2016, FCIA (Fibre Channel Industry Association).

#	FACTORS FORCING NEW APPROACHES	NATURE OF IMPACT	WHAT TO LEASE
A	100% increase in 40GbE Port counts per year	Line cards (in chassis-based switches) and stackable switches have fixed port counts. When the number of cables exceed those, they must be replaced with different ones. Also, the 40GbE technology is based on older 10G-lane technology, and is being replaced with the new 50GbE technology, based on the 25G lanes—which will force replacements in the near-term.	<ul style="list-style-type: none"> • Cards (some) • Stacks (all)
B	20% server growth per year	Servers feed Top-of-Rack switches , and typically have two network ports for redundancy. 20% more servers means 40% more cables/ports—every aggregation switch upstream must be enlarged (bandwidth, ports, security, etc.) or replaced. Servers are the largest consumers of bandwidth, and growth in servers will require growth in bandwidth (via replacement of switches and/or expansion of switches).	<ul style="list-style-type: none"> • ToR (all) • Aggr (most)
C	24% growth in traffic within data centers (2014-2019)	This traffic is between servers and other servers, requiring that much more growth in Core and Aggregation switches. Many of these switches cannot grow capacity that much without being replaced by another unit. In some cases, additional aggregation switches might be needed to allow load balancing for hotspot elimination/reduction.	<ul style="list-style-type: none"> • Core (some) • Aggr (most)
D	25% growth in traffic from data centers to users (2014-2019)	Outbound traffic from the data center to end-users will either go outside the router to the internet or WAN resources (requiring expansion of the router capacity and speed) or stay inside the router (on the local campus or in the building) which will require expansion in closet switches (edge) and expansion or replacement in upstream aggregation switches . Additionally, since much of this data will be delivered to wireless endpoints (e.g. laptops, mobile devices) the Wi-Fi infrastructure will likely need upgrading/expanding.	<ul style="list-style-type: none"> • Router(?) • Edge (most) • Aggr (many) • Wi-Fi (some)
E	30% growth in 10GbE ports per year	Top-of-Rack, End-of-Row, and stackable/modular switches have fixed port counts. When the number of cables exceed those, they must be replaced with different ones. Modular and stackables can scale-out, but only to the limit of the interconnect and/or backplanes. This limits the scale-out reach, and forces a replacement. Also, since many of these ports will be carrying traffic for uplink, upstream aggregation switches may be subject to the same expand/replace pressures.	<ul style="list-style-type: none"> • ToR (many) • EoR (all) • Stacks (many) • Aggr (most)
F	30% virtualization density increase per year	As virtualization density increases (i.e. number of virtual machines/VMs running inside a physical server), network traffic out the physical network interface card (NIC) increases also. When 5 physical servers are consolidated into one physical server (as VMs), then the network traffic from the original 5 physical servers (which would have been spread across 5 physical NICs), is consolidated onto only ONE NIC. The total bandwidth requirements do not go down (except where the 5 VM's are talking "among themselves"), but the bandwidth of the server-port, cable, and switch-port must increase by that aggregate load factor (and be very resilient also—since a fault would affect 5 servers and not just one in this scenario). This means that the physical server's NIC may need to be replaced by a faster NIC (10 or 25 or 40GB), and then the upstream EDGE switch must have a matching port, and the upstream aggregation switches must be able to grow to the bandwidth ceiling.	<ul style="list-style-type: none"> • Server (most) • Edge (most) • Aggr (many)
G	31% growth in traffic between data centers (2014-2019)	Outbound traffic from the data center to other data centers (either cloud or company data centers) will either go outside the router to the internet, to WAN resources, or to dedicated core backbone lines/switches. In many cases, a 31% increase in load will typically require a change-out of line cards (for both port-types and traffic density) in both core routers and core switches. At this level of density, latency will need to be managed closely—typically eliminating the use of over-subscription, implying expansion/replacement of large/expensive units.	<ul style="list-style-type: none"> • Cards (all) • Router(all) • Core (many)

#	FACTORS FORCING NEW APPROACHES	NATURE OF IMPACT	WHAT TO LEASE
H	45% of organizations cite bandwidth requirements as a barrier to cloud use	Several use cases for external cloud services have become accepted, some of which are either “data-heavy” or latency-sensitive (or both, as in streaming media). If a firm has decided to pursue the use of such resources, network upgrades may be required, both at initial deployment and then continuing, as cloud usage by the firm increases. Investments in network equipment should be made with a view toward replacement or expansion, in the cases when the bandwidth is not needed today—but is expected to be needed in the near future. This will certainly require larger routers , core switches, and aggregation switches, but some of the edge switches might not require change (unless the LANs are themselves a bandwidth constraint).	<ul style="list-style-type: none"> • Router (all) • Core (all) • Aggr (all) • Edge (?)
I	60% of servers in data centers still using 1GbE ports	This means that many organizations are on the cusp of a major server refresh. These servers will likely be replaced by either servers with a ‘traditional’ 10GbE port or the newer 10GbE/25GbE NICs. In the former case, most server edge (ToR) and server aggregation (EoR) switches will need replacing (chassis and/or line cards). “Closet” Edge switches facing users (as opposed to servers) and their upstream aggregation servers should not be affected. A large scale move from 1GbE servers to 10 or 25 GbE servers is a radical topology change, and much of the previous network gear would be unusable in the new architecture.	<ul style="list-style-type: none"> • ToR (all) • EoR (all) • User Edge (no) • User Aggr (no)
J	60% of US network devices are aged or obsolete	This is a major security, network management, labor utilization and growth problem, which was largely caused by budgeting for too-long lifespans. A major refresh is needed NOW, and the magnitude of this refresh will require financing to avoid a major cash/capital drain on organizations. Furthermore, to avoid being at this problem point again in 3-4-5 years, the replacement equipment should be put on a short-term lease structure—to force review and accountability at the appropriate point.	<ul style="list-style-type: none"> • ALL replacement gear
K	70% of enterprise cloud usage is for file storage	The use of an external cloud for file storage is clearly data-transfer intensive, and generally not latency-sensitive. Examples of this would be backup/recovery, collaboration, and support for mobile workers. If/as these files become media-class (e.g. training videos, audio-annotated presentations), latency may become an issue and bandwidth load may even be severely taxed. For example, simple email (including spam handling—which is 70-80% of the bandwidth load) burns about 2.7KB/sec, 80KB/sec for printing, a visit to the front page of a commercial website might be 400-1000KB/sec, SD video around 200KB/sec, and HD video at 800-1400KB/sec. This kind of traffic is bursty, and would require bandwidth upgrades mostly in the main egress devices (router, core switch).	<ul style="list-style-type: none"> • Router (all) • Core (all)
L	74% of WIFI infrastructure cannot handle mobility	74% of the WAP’s in the US have speeds below 802.11n, and hence are fairly useless for current mobility use in enterprise settings. Some firms are no doubt starting refreshes to 802.11ac (skipping 802.11n), but many are only doing the 11n standard. With the massive increase in wireless loads forecasted, any new-installs of access points should be put on a short-term lease/refresh program—to facilitate selective upgrades earlier than planned. This will be an ongoing challenge to manage, so nimbleness will be the watchword here. Upstream switches—used mostly for Wi-Fi traffic—should also be looked at carefully for that reason. The edge switches will likely require in-step refreshes, but the aggregation switches might not.	<ul style="list-style-type: none"> • WAP’s (all) • Edge (most) • Aggr (?)
M	75% of WIFI access switches cannot support 10GB uplinks	In this case, the Wi-Fi bottleneck is not at the wireless access points (WAPs), but at the switch that the WAPs all feed their data streams into. These are sometimes edge wall-mounts or closet devices, but those without 10GB uplinks will need replacing soon (with units that DO have them), and then again—fairly soon—with units that have 25GB uplink capacity. Wi-Fi traffic in the past was less media-dense (because the mobile devices couldn’t handle the processing requirements), but this has changed dramatically in only 24 months.	<ul style="list-style-type: none"> • Edge (all) • Aggr (some)

#	FACTORS FORCING NEW APPROACHES	NATURE OF IMPACT	WHAT TO LEASE
N	Big Data requires 20GB LANs	Big Data has to be located close to the compute resources that use it heavily. It simply cannot be moved around very often, due to transmission costs. So, if it is kept in the Cloud, the VMs which use it heavily must be co-located in the Cloud. In this case, there is no major hit to the enterprise network infrastructure. If, on the other hand, it is kept on-premise (e.g. due to its being captured, stored, or generated on-premise), then the servers which constantly hit that data must be closely located to it. Since Big Data is used by more than one user and by more than one server, it has to sit on high-speed LANs. The minimum speed of this LAN is 20G, according to practitioners, and this requirement will grow at the same rate as our (a) usage grows; and (b) the data volumes grow. The NICs in the user workstations will need to be at the 10G range in the next refresh, and the 25-50G range within another 2 years. The data servers will need capacity up in the 40/50/100 GB ranges, with all of the LANs probably dedicated to that single usage. This is expensive gear, and for a fast refresh will require the economics of leasing to make it even ‘tolerable’ financially! To the extent some of this data sits on a shared SAN, the SAN switches and/or fabrics will need to be able to be upgraded/expanded/replaced whenever needed—and not throttled down by depreciation cycle requirements.	<ul style="list-style-type: none"> • NICs (all) • Edge (all) • SAN (all)
O	Move to new Ethernet speeds (25/50GB, new 100GB)	Historically, Ethernet developed 40GB and 100GB technologies a long time ago, but cost factors—and heat factors—limited the adoption rate. 40GB has become popular recently, but in the meantime a better standard was developed and pushed out the doors—the 25GB single-lane technology. It has a lower cost profile, and better performance/resource balance. The 50GB version is simply 2 of the 25GB lanes, and the 100GB version is 4 of the lanes. The existing 100GB standard is 10 lanes of the old 10GB standard, and the new 100GB is 4 lanes of the new 25GB—a definite advantage. The 25GB standard was only finalized/approved in mid-2016 (along with a new set of cable standards too—Category 8), so the industry has not produced the large volume of offerings and large repository of experience (i.e. ‘scar tissue’) needed for rapid adoption. But it will be here soon—adoption is proceeding quickly—and our 40GB installed base will be sub-standard and too-expensive to run (comparatively). This means that any 40GB/100GB investments today, should be made with a view toward replacing those with 25/50/100GB within 2-3 years. Since these are new ports (though capable of 10GB legacy protocols), adoption of this will be somewhat disruptive. Some/many of our current switches will need replacing with 25/50GB-friendly models. There is still a great deal of uncertainty as to how to plan for this in a specific company, and so nimbleness is the watchword again.	<ul style="list-style-type: none"> • Edge (all) • Aggr (all) • Core (most)
P	Availability of 40GbE options in converged devices	The point made above about the 25G versus 40G technologies applies here as well. Cisco, for example, has recently released 40GB add-ons to some of their switches that were designed for ‘ converged ’ architectures (e.g. fabric switches to serve both Fibre Channel and Ethernet subnets). Business needs might require investment in this ‘older’ technology, but should be done with ‘nimbleness’ in view—fostering the 25/50 offerings as they become more compelling and more mature.	<ul style="list-style-type: none"> • Hybrid (all)
Q	Data center segmentation (DR, growth into colo's)	Oddly enough, this is the opposite trend to data center consolidation. Recent figures on data center build-outs and investments point to more use of co-location facilities by enterprises. This is sometimes due to spill-over (where the DC is maxed out in space, power, etc.), sometimes due to proximity to technology resources (e.g. fiber, data ‘supply-chain’ partners), and sometimes due to preparation for divestiture of a business unit. In any case, larger ‘chunks’ of network equipment (e.g. a giant core switch, a bank of aggregation switches, a rack of security devices) have to be broken up into smaller ‘chunks’ for segregation. This can be an expensive proposition, if the larger chunks cannot be economically ‘converted’ into smaller ones (via leasing). In a purchase scenario, a firm would buy all smaller ‘assemblies’ in preparation for this, but that does not lend itself to economies of scale that comes with the larger, more-integrated devices. The only way to reduce this risk is via lease or rental finance.	<ul style="list-style-type: none"> • Core (all) • Aggr (much) • Edge (no)

#	FACTORS FORCING NEW APPROACHES	NATURE OF IMPACT	WHAT TO LEASE
R	Inability to forecast demand growth	In many ways, this is almost a summary of all the above factors. These facts and trends set some ranges for us, but these are just not reliable enough. They warn us about uncertainty and about unplanned change and about unprecedented demands upon the IT infrastructure, but they cannot tell us what to plan FOR, nor what technology to invest IN. This uncertainty amounts now to investment risk and operational risk. The only risk mitigation tool we have here—in the area of rapid technology change—is the use of returnable assets (e.g. rental-type leasing). Various pay-as-you-go models have been advanced, but so far the economics have not been feasible to the supply-chain and end-customer.	<ul style="list-style-type: none"> • All hardware
S	Availability of SDN (software defined networks) options	Uncertainty about growth rates is one source of risk; uncertainty about future-friendly architecture is another. The rise, rapid development, early success stories, and commercial availability of the SDx-design products has many large shops rethinking their next round of investments in the core . Two of SDx technologies are relevant here: SDN (software defined networks) and SD-WAN (software defined WANs). Although the carrier and service provider industries have done the most in these areas (with NFV – Network Function Virtualization—as a form of SDN—being the major area of development by the larger telecom vendors), there are a score of commercial products available to the enterprise already. Skills sets are scarce, but the promise of separating the control plane from the data plane—allowing the use of commodity hardware instead of specialty hardware—has many firms delaying network investments while they evaluate this. This, again, is uncertainty, and investments in routers , core , and aggregation switches should be loosely held—in case these newer options become attractive.	<ul style="list-style-type: none"> • Router (all) • Core (all) • Aggr (much)
T	Availability of NVMe on Fabrics (e.g. FC-NVMe)	The availability and adoption of flash storage arrays revealed the I/O bottleneck that disk controller protocols created. The access speeds of flash/SSD drives were faster than our controllers could handle, and this led to architectural work-arounds that bypassed those controllers. One industry standard that emerged from this was Non-Volatile-Memory-express (NVMe) – a method of treating a block of flash memory as a non-disk device. This protocol could be layered on top of hardware protocols (e.g. the PCI-e bus) and the disk controller protocols (e.g. SCSI) could be layered on top of NVMe. These hybrids, though, required all of the components to be physically close together, so NVMe benefits were not available to external flash storage arrays. But the industry developed new standards in 2016 in which NVMe could be layered over fabrics, and one version of that is for Fibre Channel (FC-NVMe). These will be new adaptors for SSD/flash arrays , that layer the FC protocol over NVMe, allowing the distance-friendly FC tools to be exploited. These are different adaptors and ports, and would require replacement of older ones—when adopted. For firms needing the speed of Flash with the tools and sharing architecture of Fibre Channel, this technology would be something to invest in—although it might require decommissioning of older ports. Again, uncertainty suggests holding the assets 'lightly'—or at least some percentage of them.	<ul style="list-style-type: none"> • SAN adaptors and ports (some)