Fabric Connect / Shortest Path Bridging: The Quiet Revolution

Extreme Networks goal is to deliver communications solutions that meet the expectations of enterprises operating in the digital age. With Fabric Connect, Extreme is leveraging simplicity to create agility – making long wait times and design constraints a thing of the past. This simplification speeds application deployment, streamlines service provisioning, and enables far greater network stability and resiliency. This paper examines the advantages of Fabric Connect / Shortest Path Bridging over traditional networking approaches.

Eliminating Complexity: From 10+ Protocols to 1

When looking at conventional networks built over the last 20 years, it is easy to observe the successive layers of complexity that have accumulated. VLANs were introduced for Layer 2 virtualization and segmentation. Dynamic IP Routing protocols (RIP and OSPF) were introduced for Layer 3 forwarding. IP Multicast was introduced for broadcast applications and brought with it IGMP (at Layer 2) and PIM (at Layer 3) as new protocols. And BGP was introduced to handle route forwarding and peering to Internet providers. And this is just a snapshot of the protocols running in most enterprise networks today.

This conventional (or legacy) network architecture has clearly reached a very high level of complexity. More critically, the protocols themselves have a very high level of interdependency. For example, if a problem or failure occurs at Layer 2, then the protocols above – and related business applications – are also impacted. It can be likened to a "House of Cards" where if the stack of cards collapses, it triggers a costly business outage. This is further exacerbated by the slow and unsynchronized re-convergence of these multiple inter-dependent protocol layers.

A powerful feature of Extreme's Fabric Connect is that it makes these protocols redundant and unnecessary. In fact, when migrating from a conventional network design to one based on Fabric Connect, the following protocols are no longer required: STP, MSTP, RSTP, RIPv1, RIPv2, OSPF, EIGRP, ECMP, PIM-SM/PIM-SSM, DVMRP, LSP, MPLS.

What is equally powerful is that Fabric Connect can also run in parallel with any or all these legacy protocols. This provides a very smooth migration to a Fabric-centric architecture and ensures that any design-specific requirement, however niche or sophisticated, can be seamlessly accommodated.

Faster Time-to-Service: 11x Better with Edge-Only Provisioning

To demonstrate Fabric Connect's many benefits, independent research was commissioned to document what Fabric Connect customers were experiencing in their real-world deployments[1]. The results were eye-catching! They included a 11x improvement (or 91% reduction) in implementation time and a 7x improvement (or 85% reduction) in both configuration and troubleshooting times; and a 2,500x improvement in failover times. In fact, outages caused by human error were virtually eliminated (see Table 1).

This research reinforced an earlier report measuring the impact of network changes (or updates) on the business[2]. Most network changes necessitated a maintenance window, with IT networking professionals having to wait – on average – 31 days before they could make the changes necessary to the corporate network.

Fabric Connect's unique edge-only provisioning model is a primary contributor to these observed benefits – including the 11x faster time-to-service. Edge-only provisioning minimizes the number of network changes required while allowing them to be made in real-time. This radically improves service agility by eliminating the typical maintenance window a business must wait for a network update. Instead of waiting a month on average, a change can now typically be made the same day. And once the decision is made for the change, IT can almost immediately deliver the service to the business.

Time-to-service is further improved with Fabric Attach – an open technology that has been adopted by Open vSwitch[3] as well as a number of end point and third party switch vendors. Fabric Attach enables “auto-attach” of endpoint devices, allowing businesses to dynamically deploy services and devices in a highly dynamic environment. Fabric Attached endpoints connect to the appropriate network resources; whether if it’s a Fabric Connect Virtual Service Network (VSN), or a conventional VLAN.

So, whether it’s extending Campus Wi-Fi services or supporting dynamic Virtual Machine (VM) activation in a Data Center, Fabric Connect (complemented by Fabric Attach) is unique in simplifying network operations and improving service delivery.

Metric	Before	After	Times Better	% Less	% Seeing Improvement	Paraphrase
Implementation Time	14 Days	1.3 Days	11x	91%	68%	Weeks to Days
Configuration Time	4.6 Days	0.7 Days	7x	85%	86%	Days to Hours
Troubleshooting Time	39 Hours	6 Hours	7x	85%	41%	Days to Hours
Failover Time	>13 Minutes	0.32 Seconds	>2,500x	100%	70%	Minutes to Milliseconds
Outages (Human Error)	3/Year	0/Year	-	100%	74%	De-Risk
Wait Time	31 Days	1 Days	31x	97%	-	Month to a Day

Table 1: Fabric Connect / Before and After Network Outcomes

Better Time-to-Repair: Eliminating Hop-by-Hop Gives 6.5x Improvement

The Fabric Connect Customer Experience Report also highlighted improved service availability when outages did occur - driven by more efficient troubleshooting. Customers reported decreased troubleshoot times due to network outages, with on average a 6.5x (or 85%) improvement. The highest reported improvement was 8.5x (or 92%).

Even more interesting was the very sizable group (41%) of deployments that could not easily quantify a TTR improvement; simply because they had not yet experienced a network issue since their implementation of Fabric Connect.

The basis for these improvements can be attributed to the underlying protocol that powers Fabric Connect; IEEE 801.2Q Shortest Path Bridging (or SPB). This single, unified protocol builds a robust and resilient topology that facilitates troubleshooting.

With SPB, services are only “presented” at the edge, and then only on those nodes that specifically interface these specific services to local endpoints. Core or distribution nodes are involved only in data forwarding and never terminate services. This minimizes the fault domain that needs to debugged to a per-service basis. No longer does every node in the entire network have to be evaluated and potentially analyzed. Troubleshooting a Fabric Connect service involves the direct analysis of just two nodes, regardless of how many intervening and interconnected nodes may exist.

To add visibility at the SPB layer, IEEE 802.1ag provides advanced operational, administrative, and management capabilities. This means that the days of laborious hop-by-hop troubleshooting of the FDB, ARP, and IP Route tables is over.

Enhanced Business Continuity: 13 Minutes to 320 Milliseconds

Network availability – or more precisely, unavailability – is probably the most important metric to the business, and clearly the most visible. When the network is down, there’s no place to hide.

Again, looking at the real-world experience of our Fabric Connect deployments, three-quarters of these customers reported significant improvements, compared to their legacy network. On average, an improvement of more than 2,500x was observed, or more than 99.999%.

Recovery times went from 817 seconds to 0.32 seconds. And it can be appreciated that 817 seconds (or more than 13 minutes) is a lifetime in terms of application state-awareness and the end-user experience. Whereas, 0.32 seconds is hitless.

What's highlighted here is the knock-on effect that network outages and extended failovers have on business applications. While conventional Layer 2/Layer 3 redundancy protocols can theoretically be tuned for sub-second failover, most real-world deployments, especially in large and complex environments, come nowhere close to this. 40 second failover times (or more) are not unusual, but that is only for network recovery. In that time, applications databases have gone out of synch; end users are seeing applications failure; IP phones are rebooting due to the loss of connectivity with their call servers; and the list goes on. There is a domino effect that has a very costly impact to the business.

In contrast, Fabric Connect consistently delivers sub-second recovery; that's full network recovery for Layer 2, Layer 3, and IP Multicast. And because the network has effectively never gone away, upper layer communication protocols are totally unaffected – as are the individual end-users and the business as a whole.

Fabric Connect Customer comments:

"It has made identifying issues much easier, but there are hardly any issues these days anyway."

"Touching wood, we haven't really had any problems to troubleshoot since implementing Fabric Connect, so that number is a bit of an estimate."

"The 'since implementing' figure I have given you is a guess, as we haven't had to troubleshoot as yet."

Invisible Core: Simplified Security

Fabric Connect delivers significant differentiation, when it comes to network visibility and traffic transparency.

Enterprise networks have typically evolved into a collapsed, routed backbone with multiple virtual interfaces providing connectivity to a variety of user segments. In practice, this is configured as multiple Virtual LANs (VLANs) servicing groups of users, each with a routed interface (terminating on Virtual Routers). The Layer 3 engines at the heart of the network populate tables with all known routes, creating a situation where any-to-any connectivity is the default behavior.

This technique works but it also introduces some undesirable characteristics:

- In larger networks, end-to-end connectivity is provided by a series of hop-by-hop forwarding decisions. Configuration is often complex, especially when Layer 2 VLANs need to span beyond a single physical location. Configuration scripting can aid bulk changes but is also prone to input error.
- Being IP-centric, the conventional network topology is very easily mapped; good for network management purposes, but a double-edged sword as it provides an easy attack platform for hackers. Attacks can be launched from any point within or external to the network.
- To control the default any-to-any network behavior, businesses often lock-down connectivity to selective paths, so that any-to-any doesn't become a vehicle used by attackers. Options include using Access Control Lists (ACLs) or distributed physical or virtual firewalls. These measures can be expensive and complex to plan, deploy, and maintain.

Fabric Connect delivers a distinctly different administrative experience. Rather than any-to-any connectivity, Shortest Path Bridging uses one-to-one, or a series of multiple one-to-one, mappings. Services, or Layer 2 and Layer 3 "Virtual Service Networks" (VSNs), are a function of explicit provisioning, and communication between different services is blocked unless specifically enabled.

The Fabric Connect philosophy delivers several benefits:

- Edge-only provisioning removes any need for service-specific configuration in the core or any intermediate Fabric Connect node. This changes the configuration paradigm from hop-by-hop to end-to-end. Configuration is vastly simplified, and change is de-risked.
- Being Ethernet-centric, the Fabric Connect network topology is invisible from an IP perspective. There are no inherent hop-by-hop IP paths to trace, therefore the network topology cannot be traced using remote IP-based tools. Network management is fully supported. However, an individual host will only ever see, at most, the other hosts on their specific VSN. Individual SPB nodes are not visible from other hosts on any VSN. If enabled, ICMP only shows the VSN Edge nodes, and nothing of the inner network.
- Built natively as a series of isolated VSNs that interconnect specifically provisioned endpoints, Fabric Connect handles traffic forwarding in a fundamentally different way. Traffic belonging to a specific service is encapsulated with the appropriate header at the edge and remains isolated from every other service/traffic, and opaque to intermediate nodes. This mitigates the need for intra-network ACLs and Firewalls; VSNs are oblivious to each other, as are hosts on different VSNs, and there is no risk of traffic blurring between VLANs or seeping through generic routing tables.

Hence, Fabric Connect fully addresses the major security pitfalls that conventional networking introduces. But it goes further with these additional benefits:

- Traffic from individual endpoints is transported in full isolation from each other, delivering a true ships-in-the-night capability. We call this "stealth networking". This unique capability is complementary to specialized service overlays, such as PCI for financial transactions, or for data protection in healthcare, legal, and financial sectors.
- Deployed in concert with Network Access Control, Fabric Connect leverages authentication to create a very effective policy enforcement point; no connectivity is provided without hosts first proving their bona fides. Failed or suspect hosts are always completely isolated and can be mapped to quarantined or remediation zones.
- Fabric Attach facilitates the automatic attachment of authenticated end-point devices directly into their appropriate VSNs. Equally beneficial at the Wiring Closet and Data Center edges, Fabric Attach supports dynamic service attachment and removes the delays and risks associated with manually configuring conventional networks, yet seamlessly maintains Fabric Connect's enhanced security posture.

Fabric Connect Customer comments:

"We've not had a failure yet."

"We haven't seen one (a failover)."

"Prior to Fabric Connect we did not have a failover route, so I can't really answer the question for previously. Now it is 10 milliseconds."

"We haven't measured it since (implementing)... it was 15 seconds before, but it probably is in milliseconds."

"We had an outage that was particularly long so, to be honest, it wasn't a good benchmark."

"The figure I have given you is a guess, as we haven't had to troubleshoot as yet."

Automated IoT: Plug and Play Through Elastic Networks

With Fabric Connect, Extreme Networks has pioneered the concept of "network elasticity". An "elastic network" stretches network services to the edge, but only for the duration of a specific application session. As applications terminate, or end-point devices disconnect, the network service retracts from the edge. This elasticity has obvious benefits. It simplifies the provisioning of the ever-increasing number of network and IoT devices. It also reduces the network's exposure and attack profile. After all, you don't share your wallet openly with your cash exposed; you only produce it when specifically needed.

This capability is especially valuable for Internet of Things (IoT) devices. These are often unattended devices that need to be deployed in real-time, without the requirement for IT intervention or manual configuration. They also ideally are only allowed network access via a centralized policy engine (or policing device) that controls access in compliance with business policy.

Let's look at how this process works with Fabric Connect. Edge devices request their application-specific network assignment at start-up using existing authentication techniques, i.e., MAC- and/or RADIUS-based, 802.1X, and 802.1AB. These techniques can also be integrated with other network provisioning and policy enforcement tools. Network connectivity – e.g., VLAN, QoS, policy, or whatever is needed to deliver this service – is then dynamically extended to the edge device.

This ensuing networking session may last minutes, hours, or perhaps days. Regardless, the key point is that the service is automatically provisioned– or "spun-up" – without manual intervention or pre-configuration. And once the session terminates, the same networking configuration is automatically undone, removed from the edge access node and consigned to history.

Multicast Made Easy: 28x Scalability

In the early days of networking, multicast was a major innovation. But IP multicast configuration is notoriously complex. The technologies needed to make multicast work in a traditional network environment are complicated, involving protocol overlays that must be kept rigorously in synch with underlying network topologies. Furthermore, traditional multicast is not well suited for all broadcast applications. Video surveillance, for example, involves not just one source to multiple destinations, but multiple sources to multiple destinations.

Conventional IP multicast relies on a distribution tree built by a multicast routing protocol – typically Protocol Independent Multicast Sparse mode (or PIM-SM) – to deliver packets from the sender/source to the receivers that reside on different IP subnets. Multicast routing protocols also need to operate over an underlying unicast routing protocol, such as OSPF. This dependency commonly results in issues where packets transmitted by a sender do not reach receivers due to improper building of the multicast tree. In the case of PIM-SM, there is an additional dependency on a device called a Rendezvous Point (RP) to build the tree for a multicast group. Improper configuration of these protocols and functions can result in packet delivery issues.

The pseudo-state established by PIM-SM must remain in exact correlation with the underlying unicast routing topology. If this state is lost or becomes ambiguous, all bets are off. Any change to the network topology can adversely affect the stability of the IP multicast service. Additions, deletions, sudden outages for any reason (e.g., a faulty link, port or module) can all wreak havoc; the tree truncates and the distribution service for that length of the tree is effectively lost.

PIM-SM overlays are also very dependent on timers for the operating protocols and these timers must be fine-tuned. Mutual dependencies like these are difficult and time-consuming to troubleshoot, which means longer repair cycles and higher operational expenses.

Anyone that has been involved in deploying and maintaining large-scale multicast environments probably has the mental scars to prove it. Indeed, many have found it simply too problematic and have reverted to unicast, despite the downside of inefficient bandwidth utilization. However, IP multicast is making a come-back, often out of necessity rather than choice. Many technologies such as next-generation video surveillance, IPTV, digital signage, desktop imaging, financial applications, and some Programmable Logic Controller (PLC) implementations are reliant on multicast.

By contrast, Fabric Connect offers a scalable, reliable and efficient way of supporting IP multicast routing, without the onerous requirement of configuring, deploying, and maintaining a complex overlay such as PIM. Imagine a multicast network without Rendezvous Points, and complex configuration. Fabric Connect delivers IP multicast with the simplicity of edge-only configuration while offering vastly enhanced scale, performance, and reliability, eliminating PIM-induced headaches forever.

The unified, single-protocol technology that underpins Fabric Connect – Shortest Path Bridging – is extensible. Extreme has leveraged this extensibility to integrate support for IP multicast directly into Fabric Connect. This creates a seamless IP multicast capability, enabling the network to instantiate services on demand, whether they are one-to-many, many-to-few, or many-to-many. Sources are announced throughout the network using extensions in the IS-IS control plane (through defined TLVs). Receivers join a source group (a unique I-SID), by requesting membership using traditional IGMP.

In most conventional network designs, the practical limit is 500 Multicast Streams. Beyond this number, things get "extreme" very quickly, and it's normally deemed safer to build parallel networks to spread the load and risk. But imagine a Smart City environment where 10,000 or more IP video surveillance cameras are required; this sort of workaround is simply not appropriate.

Fabric Connect is uniquely positioned in being able to scale above and beyond other fabric or multicast alternatives. For example, in validating Fabric Connect with Pelco's IP CCTV solution we have demonstrated scalability well beyond 14,000 Streams; an incredible 28x improvement!

In Summary

The intent of this paper was to explain in a little more depth the characteristics and benefits that Fabric Connect delivers, placing these in the context of thousands of real-world deployments. Increasingly, organizations are seeking network virtualization technology technologies that provide dependable scaling and high reliability. They are also looking for solutions that actively promote simplicity and lower costs while delivering the agility they crave. Fabric Connect from Extreme Networks is that solution.

[1] Dr Cherry Taylor, Fabric Connect Customer Experience Research Report, Dynamic Markets.

[2] Dr Cherry Taylor, Network Agility Research, Dynamic Markets.

[3] 802.1Qcj Automatic Attachment to Provider Backbone Bridging Services.

Extreme™
Customer-Driven Networking

http://www.extremenetworks.com/contact