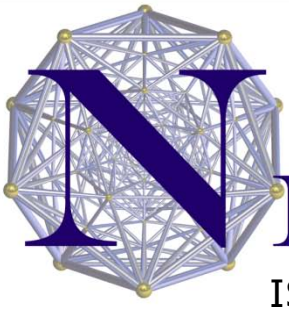

NETWATCHER

ISSN: 0890-5800

Special From our March 2011 Issue

Netwatcher is the oldest continuously published private network Technology journal in the industry. Published by CIMI Corporation Voorhees NJ USA since 1982.

Selected Content: We offer our In the Know article on Juniper's QFabric and "Chained Services" from our March 2011 issue.

Important Notices:

Netwatcher (ISSN 0890-5800) is a monthly publication of CIMI Corporation. Consult us for current pricing and see below for license class limitations. Printed subscriptions are not available. Subscriptions are accepted subject to qualification only. **This special content extract is provided only under the Netwatcher license, and may not be distributed outside your company.**

This publication is a copyrighted product of CIMI Corporation distributed on a use-license basis at a price dependent on firm size and business geography. Contact us for pricing. Subscribers in good standing are entitled to distribute this publication internally within their company freely, and to print copies and circulate them within the company. This right is effective with payment of the subscription invoice for the period in which each issue is published. Those who are invoiced for payment and do not pay have no rights of use, printing, or distribution and must destroy all copies of the material for the period involved. Once paid for, issues may be used perpetually under the rights conveyed at the time of subscription. Subscribers may not give a printed or electronic copy of this publication to anyone outside their company. Contractors or temporary employees may be given access to the publication on the premises of and under control of the subscriber, but may not take copies out of the company's facilities and control. Subscribers are responsible for any infringement that occurs as a result of the violation of our license terms by a contractor, agent, or assign.

Subscribing companies may designate up to 10 email addresses to which we will make direct distribution of issues. We recommend that one such email address represent an internal mailing list for additional distribution. We also recommend that a company library or administration process be designated a recipient. It is the responsibility of the subscriber to maintain this email list and to make changes and substitutions as necessary.

If a subscribing company is purchased by another firm, previous issues of the publication are transferred as assets of the purchased/merged firm subject to the same terms and conditions as other assets. An existing subscription is also transferred.

Subscription to this journal or to any CIMI Corporation information product does not include any implied access to consulting services from CIMI Corporation in any form. We cannot provide follow-on information, clarification, etc. on topics included in this or other publications except in the context of a paid consulting agreement.

CIMI Corporation will pay a reward of $500 US for information leading to the successful prosecution of individuals or companies who infringe on this copyright. You may contact us at the number or e-mail address below if you have what you believe to be an example of such an infringement. Contact us for details on this offer.

CIMI Corporation
1-856-753-0004
inquiries@cimicorp.com
http://www.cimicorp.com

Notes on Navigation:

This document uses Adobe Acrobat™ linkage conventions throughout. We recommend that you enable bookmarks in Acrobat to navigate easily between sections. Use the Acrobat navigation buttons or press the page keys to advance or move back in the text.

This document is designed for viewing with the ***latest version of Adobe Acrobat.*** If you experience problems reading the document and have an earlier version, we recommend you visit Adobe's website at http://www.adobe.com to download the latest version of the Acrobat Reader.

Adobe and Acrobat are registered trademarks of Adobe Corporation.

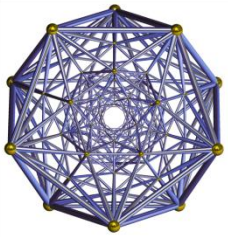

IN THE KNOW

The data center has been, for several years now, the most critical focus point for tech investment in both enterprises and service providers. For the former, data center technology holds the key to enhanced productivity and better operations, and also the key to consolidation efficiencies obtained through virtualization. For service providers, the monetization of content, the exploitation of the mobility/behavior synergy, and the realization of cloud computing service revenues all depend on the data center.

To make things more complicated, the emerging notion of a marriage of networking and IT to create a computing cloud that then supersedes both traditional network and IT architectures makes the cloud space critical in a competitive sense, and clearly the cloud is built from data centers somehow. If the cloud is important, then the "how" associated with building it is too. Because the "how" is often connected at the high level to the "why", the business case, it may be the most important thing of all.

CIMI Corporation has captured the financial importance of the data center in another way. About five years ago we noted for the first time the critical relationship between IT spending growth and GDP growth, and how periods of accelerated IT commitment coincided with the implementation of new paradigms to apply IT to business problems. That relationship created predictable sine-wave-like oscillations until 2002 when the cycles crashed at the low point and have not yet recovered their normal momentum. The conclusion that this crash and suspension of further cyclical growth was caused by the lack of a credible IT driver paradigm is inescapable.

If the data center is the heart of the cloud, then the data center network is the heart of the heart, the artery whose health determines the value of the data center by determining how effectively it can move information from storage to compute to users and back. The growth in data center IT assets has created a growth in data center networking, but unlike the IT world where guiding architectures have been offered from the first, there's been no real candidate for a unified data center network architecture that supports both the current missions and the future evolution to the cloud. Could that architecture be the key to creating the "credible IT driver paradigm" we need?

Nature, and markets, abhor a vacuum. Juniper Networks launched a program to create a new data center fabric to replace the disorderly jumble of switch layers that creates most of today's data center networks. Stratus, as it was called, was one of the first new things that CEO Kevin Johnson shared with the media and analysts when he first joined Juniper. Now, matured and productized as QFabric, this project has moved from science to reality. Does QFabric change the data center enough to create that credible IT driver?

Juniper also announced, within a week of QFabric, a new optically integrated MPLS label switch, a major step by the company into a packet optical future. Juniper claims significant benefits in capex and opex for its new product, the PTX. Our model suggests that at least some of those benefits can be proved out. The PTX is in part a response to service provider demands to increase the optical content of their IP network cores. Verizon has been particularly active in that area, and for almost a decade. The replacement of a hierarchy of big core routers by a web of active optics isn't a prospect likely to bring shivers of joy to any

router vendor, and in their positioning of the PTX Juniper is understandably focusing on the here-and-now and touting cost benefits to their PTX-based IP core versus agile optics. We don't disagree that there are benefits, but we do believe that whatever cost-based drivers exist for network procurement, they'll drive everyone to a place they don't want to be—commoditization. The value of the packet optical core, then, has to be at a higher layer. Dare we suggest a cloud?

Our view is that neither QFabric nor the PTX, taken alone, constitute a revolution because they don't fully define "the cloud". But should they be taken alone, even if Juniper's not playing up the symbiosis? The ASICs used in the two, or at least those scheduled to be used when all of the QFabric components are released, appear to be the same. The mission of determinism in network service behavior may be different in terms of the QoS parameters themselves, and the scope of the services involved, cut they're still pretty much the same. So you **could** argue that the PTX, or its concepts, can extend QFabric's value proposition. Extend it to the cloud.

We're more interested in the role that QFabric, MPLS, and PTX could play in creating that notion of distributed empowerment we opened with. The cloud starts in the data center; can QFabric be that starting point? The cloud spreads throughout the market areas. What spreads it? Could the PTX play a special role there? If so, could the PTX and QFabric be a piece of something bigger, something Juniper has (at least for now) not elected to name and brand? We think it could be, which is why we're dedicating some ink here to a specific vendor announcement. But only Juniper knows whether it **will** be.

What's the Data Center of the Future?

If you look at the "typical" enterprise in our survey, you find a company that has two primary and four secondary data centers located to serve their market geography, which is typically multi-content if not global. Primary data centers have an average of 480 servers and secondary ones an average of 131 as of the fall of 2010. These are linked to an average of 130/48 independent storage devices, and about half the total servers have internal storage as well. The network that connects the devices is on the average three tiers of switches in primaries and two in secondaries, with an average of 39 switches in the former and 13 in the latter. But averages are deceptive; the top quintile of enterprises has five times the device count per data center and nearly three times the number of data centers. Financial and health care conglomerates dominate this quintile.

Right now, server utilization in primary data centers averages 27% for CPUs, and 19% in secondaries is typical. About 78% of servers are dedicated; only 12% are virtualized taken across the whole market. The top quintile runs just a bit more utilization in primaries and 19% utilization in secondaries, and they have 18% virtualization. Clearly consolidation and virtualization have a long way to go.

Enterprises also report that they have approximately as many servers located outside of data centers as they have inside (top quintile currently has more because of branch server locations). These servers are three times as likely to be targets of consolidation, and in most cases the goal is to pull these back into data centers in some way. The top quintile of enterprises favors using private cloud or hybrid cloud technology to do this, and for that group the majority of consideration of SaaS-type office productivity services comes from these branch activities.

Among the data centers, about two-thirds of enterprises indicated that they were actively planning to consolidate within or between data centers, with the largest targets being the secondary data centers. About 15% said that they were also actively planning to use cloud technology to do this. For the primary data centers, about half were targets of virtualization; almost three-quarters of secondaries were targeted. Half the enterprises were also looking to reduce primary and secondary data center size by offloading peak loads and backup tasks to a hybrid cloud (the public portion, obviously).

In the long term, 100% of the top quintile of enterprises, 88% of the second quintile, and 65% of the third quintile indicated they believed they would have a private cloud implementation in place within 5 years. None believed that they'd be fully cloudified by the end of 2011 but 25% thought they would have made "some progress" by that time, and more than two-thirds said they would have made progress by the end of 2012.

In storage, 29% of top-quintile enterprises indicated some data centers had converged storage/network technology in place. That number drops sharply to only 12% average, though overall more than 70% of enterprises believed they would have converged SAN and LAN within 5 years. Among the top quintile the number was just over 90%. FCoE was the most-cited approach.

Converged storage network users reported that they experienced "noticeably" more problems with application performance and storage access issues under load but only 13% said that these problems were common through their operating day. Nearly two-thirds said they expected that storage traffic would be the primary driver in their changing their data center LANs, and the migration to converged storage/LAN devices was cited as the primary driver for network change by those enterprises who'd not yet converged their two networks.

SOA penetration in the enterprise varies significantly even within the quintiles by firm size. Overall, 83% of enterprises said they had "some" SOA applications, 18% said they used SOA for their primary applications. Top-tier enterprise numbers were in the same range, within limits of statistical significance. We noted that there seemed to be two groups—one who had significant SOA deployment and one who had relatively little, and so the "averages" weren't as meaningful here.

All the "SOA-for-primary-apps" enterprises said they had significant interprocess communications (IPC) traffic within their data centers, and almost a third said that some IPC traffic already existed across data center boundaries. A fifth of the SOA-committed said that IPC performance limited application performance because of network latency, and over 80% said that they expected that SOA IPC traffic would impact performance as they increased their SOA commitments. The financial industry, health care, and retail were the sectors most likely to believe that IPC traffic would create performance bottlenecks in the data center network. Enterprises reported that storage performance issues could create up to 90x variations in application performance, network performance issues up to an 80x variation, and middleware performance up to a 40x variation. Obviously nobody is letting any of these factors slow them by those maximums, but the numbers show that stuff the network can influence has a major effect on performance today, and probably would have larger effect as we tune IT up for the cloud.

If you pull all the factors together, enterprises believed that their data center networks would change "significantly" within 5 years, that they were "more likely than not" to seriously consider an alternative vendor in the data center, that flattening and upspeeding the data center network would improve performance "significantly" for their key applications, and that

latency and packet loss were "more important than capacity alone" when designing or considering changes to the data center network. All these points were accepted by the top-quintile enterprises at nearly 1.5 times the rate of the average among all enterprises.

The picture we saw, in short, was that enterprises were first committed to change in the data center and second recognized that nearly all their changes were likely to create new demands on the data center networks. However, we didn't find a significant level of understanding in how to architect an optimal data center network and we saw resistance to change (in the form of somewhat protracted estimates on when change would be complete) that likely arises from a combination of uncertainty about their movement forward and the financial and practices inertia of their current network. Enterprises were also concerned about the transition to a new architecture; 80% said that forklifts were unacceptable and 67% said that they believed they would experience unexpected problems even during a phased migration.

But what's an ideal data center network, to the buyer? Given our last paragraph, you won't be surprised when we say that they didn't have a clear picture. Only 27% noted "fabric" as a specific preference, though more than half had parametric requirements that could only be met with a very efficient flat LAN or a fabric. Fabric sensitivity was highest among users of converged storage networks. A third of enterprises indicated their future strategy in vendor/product terms, with two-thirds of this group citing a Cisco strategy and the remainder divided surprisingly evenly among players like HP, Juniper, F5, and Brocade. We do not survey buyers for current vendor commitment. Fabric alone doesn't make an ideal data center.

So the net-net here is that buyers are reconciled to data center network change, perhaps more than seeking it in fact. They don't know exactly what they would be expecting to change to, meaning that they have no specific product, vendor, or architecture goal in the majority of cases. They do have general appreciation of the need to improve performance, which included latency, packet loss, and throughput, but they tended to express their improvement goals in terms of changes from the current level ("I'd like to drop my packet loss by sixty or seventy percent") than in terms of specific numbers ("Any latency over about 3 or 4 hundred microseconds is not tolerable"). Again, companies in the top quintile and those in the financial and health care sectors in particular were likely to have more specific architecture goals and more deterministic performance objectives as well. None of the goals were strong and specific as drivers; they were more a statement of where enterprises thought they were being pushed than a justification for seeking that destination out proactively.

There is a proactive driver, one that enterprises believe embodies their specific IT evolution, and it's "the cloud". The top of the strategic issue pyramid is cloud computing, period. If a product or architecture addresses that top issue convincingly it wins a hearing with enterprise buyers. If it doesn't, then it's going to have to be sold to lower-level managers who are cost-driven and much less motivated. Build a cloud and you build a future, and to build a cloud you have to present more than the boards and nails and pipes. You need a blueprint.

Enter QFabric

To most everyone, QFabric is the productization of the Stratus concept, Juniper's long-awaited data center fabric switch. As such, it has the characteristics of a virtual switch made real, a single set of devices that doesn't so much replace **devices** as it does abstract the number of devices downward to simplify operation and normalize performance. You could end up with just as many QFabric nodes as you'd have LAN switches, but what you'd see is a single logical switch and much more deterministic performance.

QFabric is a switch architecture that creates an optically meshed interconnection of on-ramps (called "nodes"). These nodes are connected to Interconnect fabrics, and Juniper explains this as being a kind of chassis switch with the line card being attached remotely via fiber instead of being wired onto the backplane or midplane directly. A single node can have four fiber paths out of it, and each of these can be connected to an input port on an interconnect chassis with a current fiber-imposed practical limit of about 150 meters of fiber between node and interconnect. Those chassis can't themselves be interconnected.

If you visualize this, picture first two nodes (A and B) with A hosting a source port and B an egress port. A and B could be connected to the same interconnect chassis with all four of their fiber ports, in which case we'd have 4 paths' worth of capacity available between A and B, which Juniper describes as cross-sectional bandwidth. But A and B could also have their four fibers connected to four **separate** interconnect chassis. That doesn't change the cross-sectional bandwidth (because you still have four paths between A and B in effect) but now no chassis failure can take out the whole A-B connection set. A full interconnect has 22 Tbps bandwidth (44 Tbps if you measure both directions separately). The ports of a node would oversubscribe the trunks by 3:1 if fully loaded. Since any pair of nodes has to have at least one common interconnect to talk, you can't scale the fabric up endlessly unless you don't want any-any connectivity.

To manage the latency and equalize performance in this structure, Juniper makes all the forwarding decisions on the ingress port so that the structure of the fabric is invisible to the outside world at the control plane level. Packets are tagged on ingress with their handling instructions. Where services require multiple routings to successive engines, these routings are also determined on ingress. Presumably, forked or conditional routing by engines would require some manipulation of the route stack, but they didn't get into that.

The performance data Juniper supplied is impressive. Delay is in the microsecond level, with non-blocking performance and fairness at points of egress port competition. Jitter is essentially non-existent. To put the numbers into perspective, our model says that the application performance improvement that **could** be achieved by the average large-enterprise data center **for applications significantly impacted by intra-data-center network delay** is about 20 times the difference between IBM and Oracle OLTP benchmarks—the ones the two companies have been sparring over for the last year.

The problem is that data centers are, in the main, not collapsing due to complexity or growth today. Most applications don't fit the bill of being "significantly impacted by intra-data-center delay". The benefit of QFabric thus depends on something changing the data center, and the only credible "something" we have is the concept of the cloud. The benefit drivers of QFabric are compelling presuming cloud data centers, not without them. To drive cloud data centers to create a cloud is therefore the mission Juniper has to accept to make QFabric a complete success.

To create a cloud, you start by connecting the cloud data centers. We asked whether it was possible to extend this over the WAN and the answer was "Yes, as long as the connection is at Level 3", which we interpret to mean that an MPLS pipe could be used to extend a fabric beyond a local structure. Interconnecting fabrics to create a cloud is supported using IP, in other words. The question, then, is whether QFabrics connected via IP retain any black-box properties. Can Juniper take the simplification story it tells out of the single data center and into the multi-center cloud? Further, there's a question of whether the newly-announced

Juniper label switch (the PTX) is in any way capable of connecting QFabrics. Why? Because that could create a more economical and scalable metro-and-beyond cloud.

The key point here is that simplifying a data center through abstraction doesn't simplify a cloud unless the abstraction is extended upward to that cloud. QFabric is an architecture, a black box. That improves the operational efficiency of the data center by containing the multiplicity of elements inside a single abstraction. The same thing needs to be done with a cloud, so something has to extend the QFabric abstraction beyond the data center and across the globe. How that might be done is suggested by how QFabric already abstracts the data center network. QFabric is a classic black box, a virtual switch. A black box (classically) is defined by its properties since you can't know its contents. Those "properties" are relationships between inputs and outputs, which is a pretty darn good definition of what a "service" is. So what are the services of QFabric? How can we extend the QFabric concept of services to become a cloud?

The basic service is that of connection, of course. QFabric supports connection among members of an arbitrarily large number of arbitrarily large closed user groups—VPNs if you prefer. It supports, or will support, the popular Ethernet speeds and storage network connections. But Juniper says that QFabric also supports "services" beyond connection. These services are created by the routing of traffic to intermediary points—a "transit list" in a way. Each of these intermediate points or "engines" can do something to the traffic. For those familiar with SOA and SOAP, the concept is something like that of a SOAP Intermediary. Engines are essentially attached to ports on nodes, though you could in theory build a QFabric node that had one or more integral node/engines. In fact, we think Juniper will do that and so will some partners.

Security is the only service of QFabric announced so far, but it's a good indicator of what Juniper intends. Juniper's SRX products become the service engines, and link to QFabric where they can be made intermediaries on any specified flow, and thus add security to the flows as needed. In addition, the security policies and management link to the virtual machines in the data center (and thus in the cloud) through the vGW product obtained from the Altor acquisition. QFabric effectively virtualizes the SRXs to share them more effectively across hosts than would be done the old dedicated way, and these virtual security resources are then bound to virtual-machine security to create truly virtual and cloud-ready security mechanisms.

This is what we think the overall model for cloud services in QFabric will be. You have engines that map service features onto flows in a virtual way, and you have agents in devices that let their behavior be managed as components of a virtual network or cloud.

But wait! If services are created by the distributed binding of process engines with VPNs, aren't we defining a cloud and not just components of one? Presuming that we had a flexible definition of what an engine could do and a flexible way to create our transit list (including a branching, forking, and list-of-list capability that Juniper says is in there), have we not defined a kind of distributed software component binding capability? Or maybe a distributed multi-core processor thread distribution capability, or maybe a grid? I can have a service of security, says Juniper. Can I not then have a service of "answer my Jeopardy question, Watson?" Can I not represent enterprise IT, service provider services and features, and consumer behavior support as QFservices?

So there's the key point. What Juniper has defined in the creation of the highest level of QFabric abstraction is the first technical architecture that implements the cloud of the future inside the black box of the present. Since there's no limit on what "services" and "engine"

could offer, it could offer computing. That means that the conceptual service model of QFabric is in fact the model of a cloud. What the **implementation** does is set the boundaries on that abstraction in the real world. If those boundaries are not excessive, then QFabric can represent a cloud. So now we have to look at what those limitations set by the implementation might be.

From Fabric to Domain

OK. QFabric is a switching architecture that provides incredibly low latency and high performance within a data center, and that can be extended across multiple data centers through IP interconnection. So now we need to circle back to the abstraction and restate it more in QFabric terms. Here, since we're treading beyond the announcement, we'll have to use some of our own terminology.

Let's agree to call a high-level abstraction of a network-centric cloud, the very top of the fractal pyramid, a "domain". Let's then agree to call an individual data center QFabric complex a "fabric". In a simple single-data-center implementation, the two map to the same structure. If I want to extend the high-level abstraction over multiple data centers, connected by IP as we noted above, we would be sewing fabrics together to create a domain that spanned multiple data centers.

Here, we have to step outside QFabric. Juniper says that there is no provision to extend the node-to-interconnect links, so you can't build a WAN bridge "inside" the fabric. If we could, the QFabric abstraction could truly and automatically cover the cloud. If we can't do that, and if we're going to make a domain into a black box, we'd then have to do that by making a series of fabrics act like a single box. We have to extend the per-fabric abstraction to cover multiple fabrics—the domain. Is that possible? To answer that we have to look at the two network issues—control plane and management—and the application-layer issues.

At present, there is no means of unifying multiple QFabrics at the control plane level, nor any commitment to do that at the management plane. Multiple fabrics have to be multiple network locations, and multiple virtual devices. But the issues this creates at the control plane level aren't significant given that's what multiple data centers are today. The problem would arise if the multiplicity of devices impacted operations. That's a management plane issue.

At the management plane, the QFabric process is represented as a unified abstraction by a Director device, which was also not announced in detail, and by Junos Space. The role of the latter wasn't brought out in the press/analyst release but it was confirmed by the company. The whole idea is to create a single-device management model that simplifies operation of a fabric and, presuming there was a cloud extension created, the virtual fabric as a whole. This is consistent with how Juniper positioned its 3-2-1 data center strategy; virtual fabrics and virtual chassis were part of that process too. Thus, Space seems to be able to work with fabric management, and it should be possible to make Space unify fabric management into domain management.

So where we are is this. QFabric is not, in itself, a cloud structure. It's an element in one, and to create a unified structure you have to move above QFabric and add other elements. Logically, since QFabric is not a label edge router in MPLS terms, I can't hop directly onto an MPLS device. I'd need to pass through something that provided LER functionality. In theory, Juniper's analyst briefing noted, you could make a "node" any kind of service engine. Might that include LER functionality? It would seem, but nothing was said about that. Thus, we have

to look to management abstraction answers from "up above" as an old '60s romantic tune ("Puppy Love" for those interested) suggested. In this case, that means Space.

The Junos Space integration is critical to the credibility of the cloud-abstraction position for QFabric. We believe Space to be Juniper's greatest single asset, because it's the closest thing to a completely articulated binding between network and IT that anyone has announced (it's also the most under-played and under-explained by Juniper). QFabric makes an exquisitely capable platform. That platform can be extended into the network layer of a cloud through the use of Space-created abstractions. The network part of the cloud, operationally, could be a "Space virtual switch".

Space could also be used to create "services" for the fabrics that could supplement security service at the network layer, and that could couple applications at the service layer down into the fabric. Space-managed fabrics could then abstract into clouds. For the service provider, that role is critical if they're to link their investment to monetization opportunities in the content and mobile spaces. For enterprises, the linkage is critical for creating hybrid clouds—for federation-as-a-service in particular. The latter is also critical to network operators. We don't know what the specific Space bindings to either QFabric or to any supported IP/MPLS interconnect elements might be, and since Juniper didn't include any Space features in the security announcement, we can't use that to model how Space might manage fabric services.

But returning to our fabric-to-domain thing, it's less than clear to us at this point is how QFabric gets interconnected to create a cloud. Yes, the features are compelling in a single data center for many enterprises, but they're more transformational in terms of IT strategy if the features are syndicated across all data centers. We noted earlier that Juniper had told us that QFabrics could be linked via IP. Juniper also announced a new PTX switch, their first true hybrid optical core device. The PTX is, as we've noted, a label switch and there's a label switch function inside QFabric as well. Might PTXs be used to connect QFabrics? The value would be that the PTX is a high-density packet-optical hybrid that could create killer metro clouds where operator fiber was available, and the cost in capex and opex would be lower than the cost of using a traditional Ethernet or IP network to build the domain.

Label switches have low latency and low cost relative to standard routers. They're theoretically capable of more deterministic behavior. If Juniper were to use a cross-section bandwidth management strategy inside the PTX that was similar to that of QFabric, it could in theory establish highly deterministic QoS inside a PTX network. Combine that with highly deterministic data center network behavior from QFabric and you have a pretty solid network foundation for a cloud. Add in a way to bind computing to the network and you have the cloud itself.

Forget the OTN-competitive concept for a moment; forget even the "core". A PTX is a lower-cost label switch capable of supporting highly deterministic traffic. That's a mission very likely to be found inside even metro networks (the whole MPLS-TP versus T-MPLS versus PBT wars are fought over that mission). PTXs would make a better metro core or content core than a standard router network would, and PTX-based metro cores connected by a PTX national or global network would similarly be very deterministic.

But…remember that Juniper is very clear in saying that the PTX is an LSR, a label switch router. It's not an LER, and neither is QFabric. That means that as things stand, you can't hybridize PTXs and QFabrics to make a cloud. There would have to be an intermediary box (a cost element) or you'd have to add LER functionality either to a QFabric node or to the PTX. Neither is an acknowledged Juniper plan, but operator interest in this area is high. We think

that MPLS extension of QFabrics is a given, and that PTX extension is likely, simply because that's what the buyer wants and the process shouldn't be rocket science given Juniper's comments about flexible nodal "engines" on QFabric.

But remember that the cloud is an IT structure as well as a network structure. We saw that in theory QFabric could present cloud services as its own services. Can MPLS spread that representation cloud-wide, and thus make the combination of QFabric and MPLS a true cloud infrastructure? If so, then the PTX might be able to play a truly critical and symbiotic role.

That's important for the PTX as well as for QFabric, because the value of the "packet" part of packet-optical cores is in the value of connectivity and traffic optimization, both of which are increased radically by the presumption that the "core" is linking cloud components and not just hordes of zero-marginal-revenue Internet users. You can make a case for the PTX as a defense against pure agile optics in the Internet core, but you can make a compelling case for the PTX as the core of a cloud network.

The "Symbiotic Cloud" and Express Domains

Network operators like architectures because they guide how products are assembled to create specific services that monetize specific opportunities. We've used the term "Symbiotic Cloud" as the name of the cloud-computing-based, network-coupled, successor to the old PSTN Advanced Intelligent Network (AIN) architecture. Just as the PSTN dictated premises voice, enterprise voice, the Symbiotic Cloud will dictate private clouds. For QFabric to be transformational at the provider **and** consumer level, it has to not only connect a Symbiotic Cloud, it has to create one that spans both those levels—because no single architecture now does that.

A Symbiotic Cloud is a combination of network and IT elements that form a virtual computer that hosts service features, runs management applications, and offers cloud computing as a service to customers. You can figure out how to build one today; in fact, a thousand ways. That's the problem. We could build have built the PSTN a thousand ways too, and absent a controlling architecture we'd have build one of each and nothing would have inter-operated. The value of the Symbiotic Cloud notion is that if we could abstract the whole of service provider infrastructure into a single virtual computer, and then define the rules for assembling the pieces internally based on a common architectural goal, we'd have the blueprint of the network of the future. So the test of transformation potential for QFabric, and also for the PTX, is whether they can create that single virtual computer and host the Symbiotic Cloud.

One question for the Express Domain is exactly how deterministic an MPLS network could be. How does Juniper manage resources inside it, and what range of QoS might it offer? Since Juniper proposes that operators migrate their TDM traffic off separate lambdas in a DWDM fiber network and onto the label-switched paths via the PTX as TDM declines, it's pretty clear that you can create highly deterministic behavior with PTXs. More deterministic than cloud computing is likely to need, in fact. We can thus say that the first test of Express Domain suitability for the cloud mission is met; the network can create an SLA.

You see now that our abstract concept of a domain is actually important, because it's the properties of the domain that determine the properties of the cloud. Effectively we have a set of QFabric black boxes inside a bigger Express Domain black box, and it's this outside box the world sees. That Express Domain box can offer services of transport and connection that are

highly deterministic, and it can also offer (through the QFabric features) services like security that are hosted on "engines" on the QFabric at strategic data centers.

But it's still a network, and its services are still network services. We need more. In our presentation to the Street back in 2007 (the key chart of which is reproduced below) we noted that the next cycle could be expected to be called "Distributed Empowerment". Does that sound like cloud computing to you? We have believed (from our first moment of recognition of the critical value of successive paradigms to drive growth) that what's now called "the cloud" would be the next driver in our chain of technology cycles. But you can't drive infrastructure spending with a slogan, you have to produce a tangible architecture, a model to build on. That model includes the network, of course, but it also includes the computing elements, the applications, the storage, and in more than just a connecting sense. If the cloud abstracts the network, it has to abstract IT **on** the network too. With the right interpretation of features, it may be that Juniper has that model, or at least an instantiation of such a model that could ignite other attempts to create one and broader acceptance of distributed empowerment.

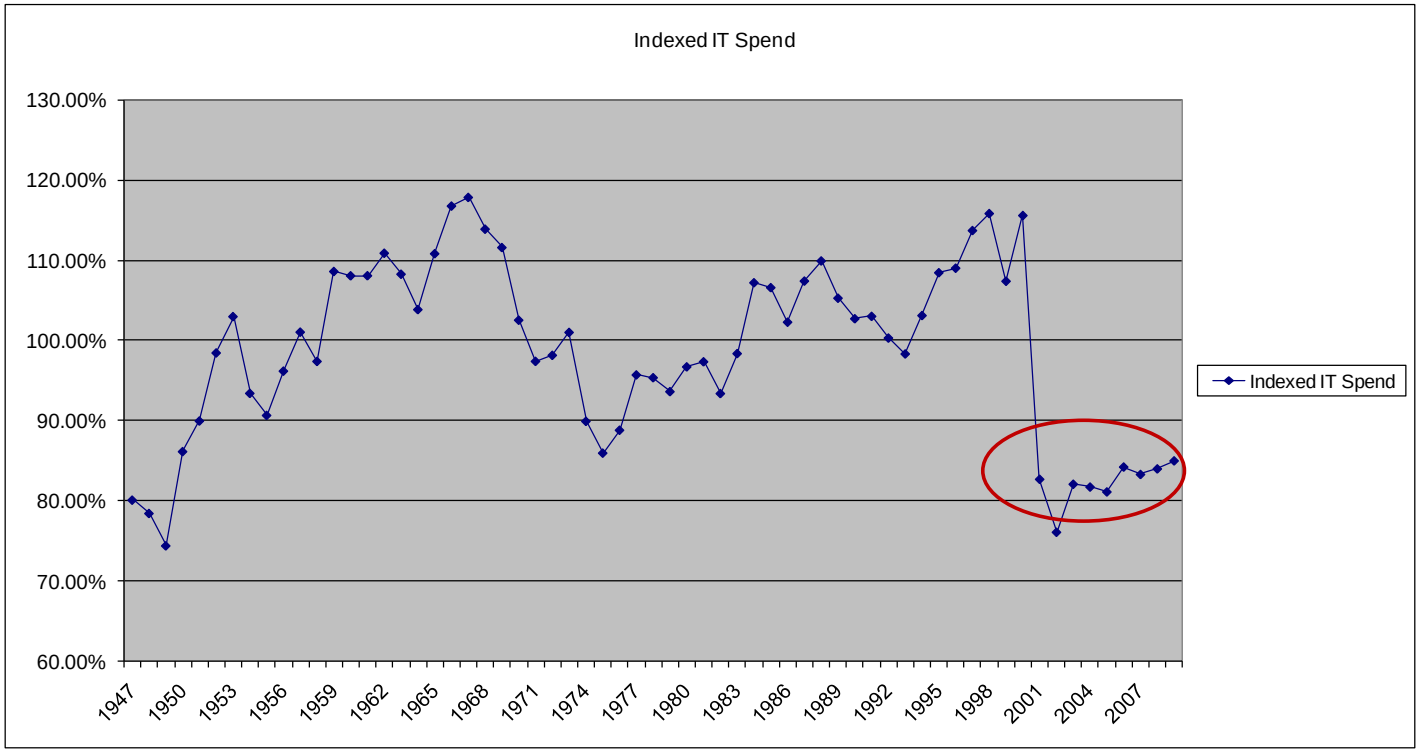

US IT spending growth relative to US GDP growth, normalized to long-term average (100%)
Source: BEA data analyzed by CIMI Corporation. Use with attribution only.

If we had an executable vision of the cloud that could be represented as a single virtual computer, we think that vision could galvanize the notion of distributed empowerment. To make Express Domains into that vision, we'd need to first have it create an explicit abstraction of the connectivity layer of the cloud, which would mean that all cloud resources could be visualized as a single virtual fabric according to the model we've developed here. Then we'd have to expand that abstraction to embrace the computing devices that are "on" the network. Neither QFabric nor any LER/LSR the PTX are computing devices, nor can they do anything direct to abstract computing devices. That means that an Express Domain has to be built at what we've called the "service layer" and for Juniper that means we're back to guess what? Junos Space.

QFabric as announced recognizes virtual networks and services. If computing is considered a "service" of QFabric just like any more traditional network service (Juniper used security as an example), then QFabric can represent the cloud within the data center. Space then has to represent it between data centers by marshalling WAN resources (like those created by the PTX) to build the deterministic underlayment of the cloud—the network—and then to couple IT assets to it to make a service layer.

We suspect that the coupling between QFabric as an abstraction and computing resources as an abstraction would come via what popular terminology would label a "Space Connector", meaning a federation bridge between a management and control input to the computing services and the Space application framework. We're told that Junos Space can represent external resources via metadata and can then make those resources available for service composition (including the composition of what QFabric calls "services"). Space can build (in abstract terms) the compute layer of the Express Domain, in short, and complete the cloud abstraction. While it appears on the surface that they could, and in fact appears that way to our fairly detailed analysis, based on what we know of Space, not only has Juniper **not** said that they can build a cloud abstraction with space, they've **not** said they could build a cloud **network** abstraction with it either. So do they now jump out and do that?

Maybe. The problem Juniper will have here is that Space wasn't positioned this way, and still isn't understood this way even by a majority of Juniper customers. Given that neither the PTX nor QFabric are positioned explicitly by Juniper as elements in a proto-cloud-element we called an Express Domain, we're building sand on top of sand by adding Space support to the Express Domain element and calling it a cloud solution. We may see things this way, but unless Juniper not only sees this but supports it, then all of this is untrue for all practical purposes and both QFabric and the PTX are simply good point product sets that solve unrelated problems. Space, QFabric, and the PTX are all announced at this point, and without any over-arching Express Domain or cloud positioning. That forces Juniper to create such a position retrospectively (which the media hates and tends not to cover, and which is therefore hard to do even at the sales level) or to risk loss of the higher-level benefit case we've described.

That may open some running room for competitors. F5 announced a reference architecture it developed for cloud computing along with IBM (also a Juniper OEM partner), but that was right after the QFabric launch and F5 didn't take any strong counter-punches at Juniper. Alcatel-Lucent, Brocade, Cisco, and HP haven't done any recent cloud announcements, so any of that group could in theory jump out with something exciting. But of course that presupposes that they have something, and that they can break with the industry trend of being boring when announcing stuff you think is exciting.

Where Does This Leave Us?

Taken by themselves, both QFabric and the PTX are significant engineering and technology advances, aimed at specific issues. As mechanisms for managing overall cost and complexity, they're strong. We even agree with Juniper's positioning of QFabric as a potential accelerant of computing performance, a step toward our utopian "Distributed Empowerment". We believe that PTX-based MPLS end-to-end services could link fabrics into a domain and create a cloud. But taken by themselves, QFabrics and PTXs are neither safe from competition nor a full solution to the problem of IT spending stagnation. Taken without Junos Space, there's no end to end context, no "Express Domain". The question we asked earlier here is critical: "Why

take them by themselves?" Why are two product families announced a week apart not an ecosystem?

We're not confused about the combinatory potential. We're pretty sure that while we may be wrong in some details here, we're right in the overall thrust. All the new stuff really does work pretty much like this, and it works with Space as well. We've had enough conversations on the details to be comfortable that we're painting an accurate picture of what could be done. There is, in an empirical sense at least, an "Express Domain".

But we're not confident either, because you can't buy an empirical sense; it has to be made real. **Look at the IT spending chart on page 10 and you understand the stakes**. The red oval shows where we have been in the last decade, a place where IT spending growth has lagged historical averages when cyclical history says it should have been leading it. Our model says that a full-on IT cycle, the kind we used to have, would create worldwide over $140 billion in incremental IT capital spending in the peak year, and over half a trillion over the cycle. If you like growth in total addressable market, how's that sound? And it's those incremental dollars, applied to business and societal problems, that can change the world—if people can apply them as an ecosystem and not as atomic products that can't individually seize the benefits. Those applications that are "significantly impacted by intra-data-center delay" in our previous reference are the applications built to exploit virtualization, SOA, and the cloud. The architecture to create them is the essential ingredient in empowerment; only after they are created can they be accelerated.

Since our presentation in 2007, nobody has stepped up to claim that they could do what we called "Distributed Empowerment" and passed the sniff test. For Juniper, the problem isn't that they couldn't back up the Claim of the Cloud, but that they're not making the claim to start off with. Could they assert that claim? Yes, in our view. That begs the question of why they don't. Do they not value the mission, or not see it? If so, they're stumbling off the path to the realization of the potential for data center fabrics and clouds to multiply the value of IT, a potential they call out in results but not in execution. The specific IT mission must be seen and realized as an implementable architecture for the cloud, because before IT can change the world in ways Juniper has outlined, somebody has to change IT.

Will network operators perhaps employ Juniper's new assets for the mission of cloud-hosting whether the company claims that mission for its own? Possibly. Will the company eventually sing its own song in perfect harmony? Possibly. We can be—we are—confident that the potential is real. As a financial analyst said, we've seen Juniper's engineering excellence emerge. The problem is that you've got to sell a great car, not just build the pieces of it and expect the buyers to do the assembly. Buyers need the guidance, and the inspiration, to bet on a different future for IT, and by betting make that future real. Sellers have to provide it. Juniper could have, and they still can, but now they can also be scooped by someone else.

Which takes us back to competitors. At the least, we expect that Alcatel-Lucent and Cisco will be doing something explicit for cloud architecture; they pretty much have to. We'd love to do a similar kind of piece on Alcatel-Lucent, or on Cisco, or on any of the other half-dozen vendors we did a service-layer review on a couple years ago. Cisco in particular has a data center product set that lends itself to cloud positioning. Similarly, we'd love to do a cloud-creation piece coming at it from the IT side. IBM's CSP2 demonstrates that you can abstract IT assets and make them into a service. Can an IT architecture abstract the network? Probably, so let's give it a try, someone; so far nobody has offered enough detail for us to be sure. Right now we don't have enough to work with, and if that changes we'll take a shot. **Offer us enough material to support your strategy, and we'll publish your approach, dear vendors!**

That's a promise. And it's one someone may well call us on, because as we've said, most of the competitors probably have some assets in the space, they're probably all developing more, and they now have a target to shoot at, at least if they look closely at Juniper's actions. We've quoted this axiom more than once: "Judge an enemy by their capabilities not by their intentions". That means competitors need to step up too, and they probably will.

We have to come back to an analogy we've made many times before. A great tire is useless except when mounted on a great car. In fact, great tires are nothing absent the buyer driver of auto purchasing. The value proposition for technology has to be seized at the highest level of consumption and then percolated downward to its execution at all levels of production. It's the top of the architectural pyramid that touches the buyer's business case. If you solve a production problem at the bottom and present no way to link the solution to the problems and opportunities the ultimate buyer has, then you've not created a revolution. In fact, you've not even addressed a secure market. Enterprises and service providers today have those high-level problems, exposing high-level value propositions. They're starting to act on them, and it's that behavior that will bring about market changes. Those changes are what will really create opportunities, and addressing them credibly requires an elevated story that's an output of elevated thinking.

A QFabric data center is credible as a cloud data center to the extent that Juniper can build a cloud from a set of them. A PTX core (metro or otherwise) is credible as more than an economizing play on the slide down the revenue-per-bit chute to the extent that it can be a cloud-builder too. Given a cloud end-game, Juniper has some great opening moves. If we're right in our assessment, Juniper has a strong technical foundation for the cloud. But what's the cloud end-game? What Juniper lacks, or has not presented, is the big picture. Will Juniper make the final connections, or will someone else? Good starts count only insofar as they create good finishes.

