ActualTech Media

Pluribus NETWORKS

Networking for the Distributed Cloud Era

Ed Tittel

CONTENTS

INTRODUCTION

Many enterprises and service providers are leveraging cloud-based services and architectures, including public cloud, private cloud, hybrid cloud, and multi-cloud. At the same time, while many workloads are moving into public clouds, many others remain on-premises in private clouds or traditional virtualized data center environments. One thing all these cloud implementations share in common is a heavily centralized architecture. This means storage, compute, and networking resources reside in one or a few centralized locations.

However, there's an emerging class of applications that have a new set of requirements that cannot be met by this centralized cloud architecture. Latency-sensitive applications such as virtual reality (VR) and augmented reality (AR) are excellent cases in point; so are data-intensive, widely distributed applications such as video surveillance and the Internet of Things (IoT). The requirements for this class of applications require resources deployed at the network edge, closer to users and things, so they can deliver on key requirements of latency, bandwidth cost reduction, autonomy, and data privacy.

UNDERSTANDING THE DISTRIBUTED CLOUD

The distributed cloud is a dynamic and flexible collection of computing resources—including networking, compute, and storage—available to users and things at multiple edge locations via the network. Distributed cloud is perhaps best defined as edge computing fused with the cloud consumption model.

Edge computing, on the one hand, pushes network, compute, and storage resources as close to users and things as is required and economically feasible. Cloud consumption, on the other hand, refers to the classic benefits of the cloud, including trading CapEx for OpEx, standing up and using applications and services as needed, scaling up and down to meet variable demand, and so forth.

As previously mentioned, the drivers of distributed cloud include a variety of factors:

- **Latency.** Many emerging applications demand low latency. AR and VR, for example, target round-trip latencies at around 7 ms, and don't work well when latencies exceed 20 ms, because it tends to induce nausea. This means putting the compute function physically close enough to the VR device to stay in that latency range.
- **Data thinning.** The cost of bandwidth drives a need to process data at the edge. For example, filtering work that teases out five minutes of valuable or informative surveillance footage from hundreds to thousands of hours of "nothing happening" need not consume excessive and expensive bandwidth. Processing at the edge makes this more resource efficient and cost effective.
- **Autonomy.** Typical IoT scenarios may involve hundreds or even thousands of sensors and actuators working together to provide, for example, traffic safety. These endpoints need to work together, even if the connection to a distant cloud is lost. Thus, processing and handling at the edge is more reliable than pushing command/control into a more distant data center.
- **Data privacy and sovereignty.** Some locales require data to be housed in the same country where users reside. The General Data Protection Regulation (GDPR) also requires organizations that store user data to be ready to show, update, and delete that data immediately upon user request or demand. All this also argues convincingly for housing and managing such data at the network edge.

DEVELOPMENT OF THE DISTRIBUTED CLOUD

Distributed cloud will have many edge locations; workload placement will be a function of meeting the application's requirements at the lowest cost. Today, the edge resides in regional data centers and colocation facilities, but this will move out into the network closer to users and things. There will be new edge locations established at central offices, head ends, and emerging edge colocation provider

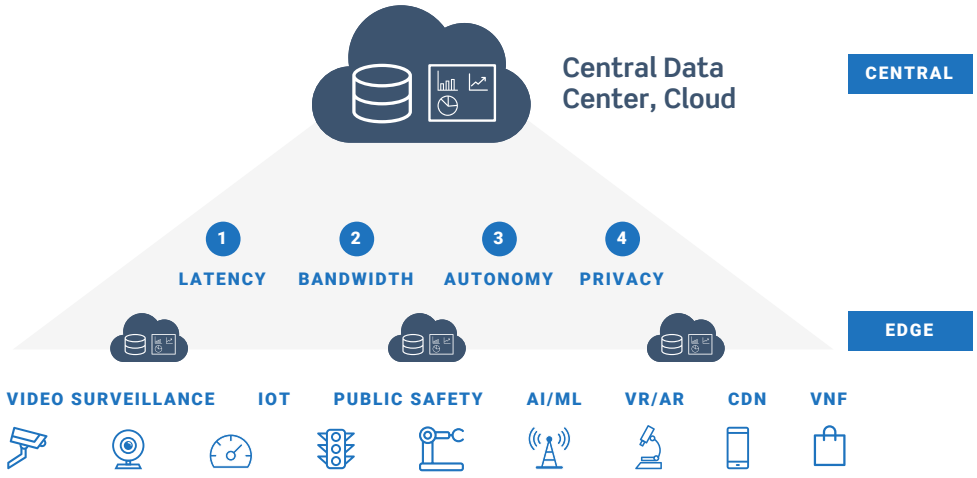

Figure 1: Distributed cloud is the fusion of edge compute and the cloud service model, which processes data closer to users and things.

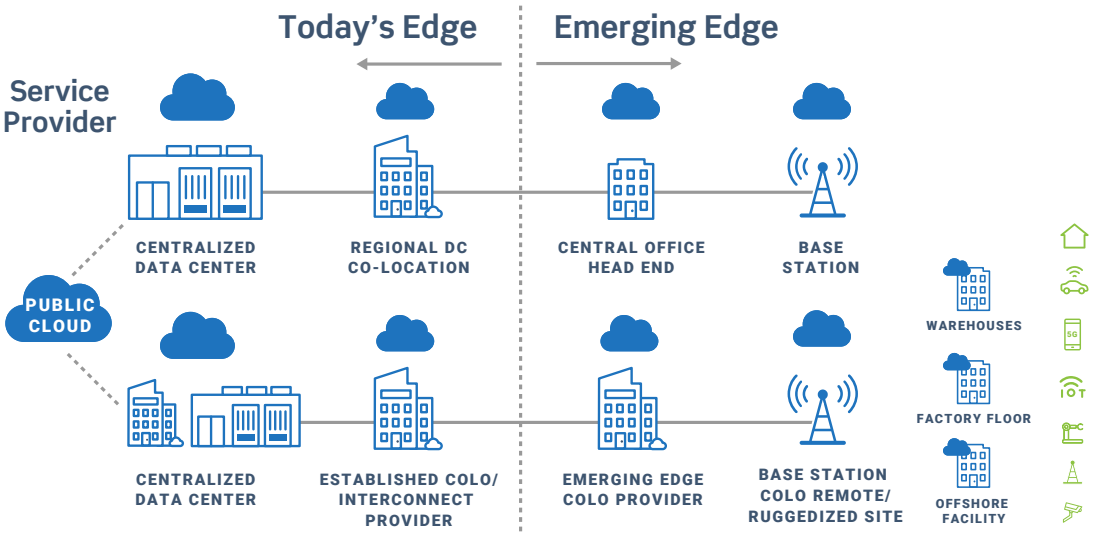

Figure 2: Distributed cloud is a decentralized architecture that will have many edge locations.

facilities, as well as on the factory floor, in remote locations for energy applications, and more. There may also be edge locations that arise at base station aggregation sites or even to the base station locations themselves as 5G increases its deployment footprint.

The emergence of 5G will be an accelerant to distributed cloud, providing faster bandwidth, reduced latencies, mandatory network slicing, and other architectural benefits to support distributed cloud. However, distributed cloud doesn't require 5G to come to fruition. The buildout of distributed cloud is already underway using multiple access technologies including 4G wireless.

To enable the distributed cloud, a number of technologies and services are being developed. This includes lower-cost and lower-power servers, and storage that works in remote locations. Some of these may need to be ruggedized for harsh environments or network equipment-building system (NEBS) compliant solutions, which are required by law in central office environments, for example.

Edge workload placement orchestration software systems are being built that will optimize the placement of workloads at the right edge location: spinning up containers, activating the application for a certain period of time, then retiring the application and containers when no longer needed.

There will also need to be systems to monitor and manage remote compute and storage, and services in place to replace failing components. And of course, distributed cloud is inherently dependent on a network: One of the most challenging aspects of distributed cloud will be how to manage networking systems that are an order of magnitude more complex.

NETWORKING REQUIREMENTS FOR DISTRIBUTED CLOUD

With a potential explosion of edge sites comes a proportional increase in complexity and cost. It's already challenging enough today for data center operators to manage servers, storage, and networking in the context of one or two centralized data centers where skilled technicians are on-site. As that increases to tens, hundreds or even thousands of remote mini- and micro-data centers, the networking challenge is daunting. In addition to operational complexity, the sheer cost of deploying so many networking devices at so many sites can be high.

An ideal solution leverages whitebox networking and provides a virtualized overlay network fabric that can be controlled as a software-defined network (SDN), and can unify multiple sites so they can be managed as one logical site.

For example, if each mini-data center site has two top-of-rack switches and there are 10 sites, it would be ideal to create a fabric that makes those 10 sites (20 switches in this example) look like one logical unit, effectively countering the complexity increase due to the greater number of remote edge sites.

The fabric should be programmable and automated from a single point of control via CLI or API, and ensure homogeneous configuration changes across all switches in the fabric. Of course, in edge scenarios, the SDN control must be implemented so there's no latency incurred or risk of communications channel failure when connecting to the controller for network state or requesting network services.

The fabric must also be sliceable (called "segmentation" in enterprise parlance) across the management, control, and data planes to support multi-tenancy and full application isolation. One example of this would be segmenting IoT data and its attack surface from valuable corporate network traffic, or providing different performance characteristics for different applications.

Ideally, the fabric implementation is developed so that it can easily work in brownfield environments, such as with existing spine switches, data center gateway devices, and WAN transport equipment from traditionally vertically integrated vendors. And, of course, the network must be highly resilient, with no single point of failure.

From the enterprise perspective, or the perspective of a managed service provider managing the network and workloads on behalf of the enterprise, there may be multiple colocation facility providers involved. Each of those facilities may or may not have their own SDN-controlled, cross-connect fabric. However, the enterprise will want a complete view and control of its own fabric, regardless of which colocation facilities are used. Thus, there very well may be overlapping SDN fabrics in this new, distributed cloud world.

Similarly, telco and cable service providers, along with cloud service providers, will need a similar capability as they build out telco cloud infrastructure for network function virtualization (NFV) where the functions might be distributed in a large number of central offices or head ends.

NETWORKING SOLUTIONS FOR DISTRIBUTED CLOUD

There are three fundamental approaches that can be used to deploy a network fabric for distributed cloud.

- **Traditional data center switching.** These offerings typically come from established vendors and are closed systems where the hardware and software are vertically integrated and purchased from a single vendor. These solutions can provide a fabric, but still require complex box-by-box configurations for configuring, provisioning, and troubleshooting.
- **Controller-based SDN fabric.** This brings automation benefits via an external SDN controller. However,

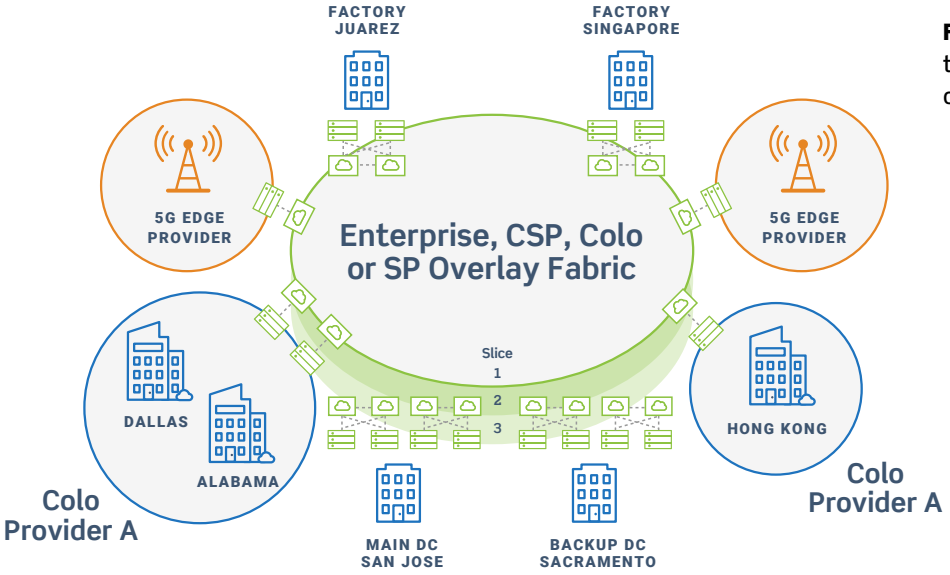

Figure 3: Adaptive Cloud Fabric networks together multiple geographically distributed sites into one logical fabric.

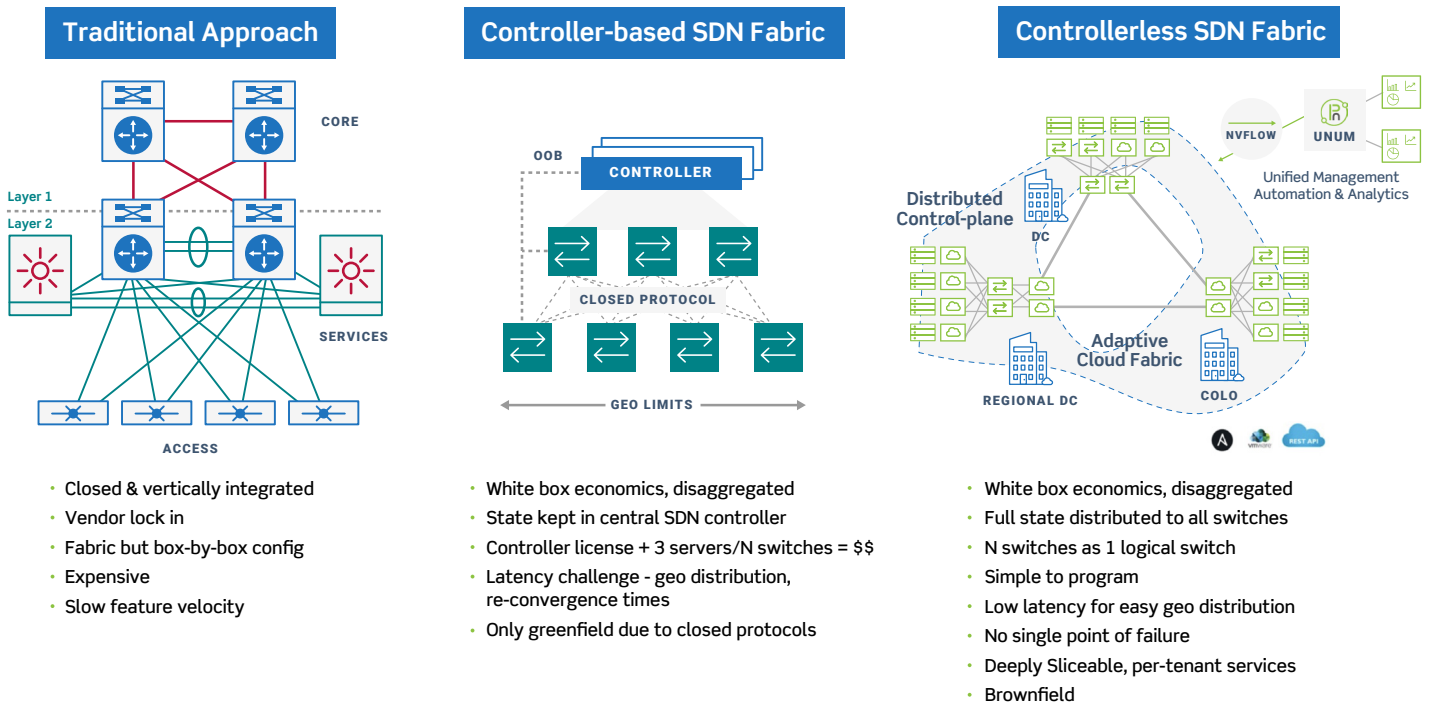

Figure 4: Controllerless SDN architecture is ideal for distributed cloud deployments.

there are challenges with this approach for distributed cloud deployments with multiple geo-distributed sites. They include the cost of the controller license and the underlying servers for every n switches; latency incurred when communicating back to a central controller; a single point of failure in the out-of-band (OOB) management network; a single management port on the switches; and challenges interoperating with existing installed infrastructure.

- **Controllerless-based SDN fabric.** This approach provides the automation benefits of SDN control, without the drawbacks. There's a centralized view of the state of the network; that full state, along with the SDN control intelligence, is distributed to all switches in the fabric. Network services are also distributed to every switch, which results in minimal latency to the SDN control function for new flows or reconvergence events. Furthermore, this approach uses in-band control communication, providing multiple paths between every node and multiple ports on the switch, eliminating single points of failure.

PLURIBUS NETWORKS SOLUTIONS FOR DISTRIBUTED CLOUD OVERVIEW

Pluribus Networks' Netvisor ONE Network Operating System and its Adaptive Cloud Fabric (ACF) deliver a controllerless SDN fabric that uses standard VXLAN overlay for network virtualization, with distributed services. It enables a large number of switches to be managed as a single logical entity.

The Linux-based, open source FRRouting-based software runs on bare metal leaf and spine switches that are Open Compute Project (OCP) and Open Network Install Environment (ONIE) hardware-compliant, providing an open, secure, and programmable next-generation network OS. This includes hardware from well-known vendors and suppliers such as Dell EMC, D-Link Systems, and Edgecore.

These open, disaggregated hardware and software solutions avoid vendor lock-in and allow purchasing departments to focus on pricing, delivery, and other key contract terms to ensure the lowest cost of capital; this is particularly important in multi-site distributed cloud deployments where costs are an important selection criteria.

Using ACF also delivers lower operational expenses owing to reduced complexity and better performance. The single pane of glass and underlying infrastructure that Netvisor ONE and ACF supports by interlinking and sharing network and configuration data ties all switches together. This means they look, act, and may be managed as one single cohesive network infrastructure, irrespective of location or device specifics. The entire fabric can be controlled from any switching node in the fabric, whether via CLI, a single REST API or the UNUM GUI management solution.

The system is also designed so that configuration changes using a fabric-wide object are automatically committed to all switches in the fabric, and if one of the nodes cannot execute the new config for some reason the commit does not take effect on any switch. This level of automation ensures consistency and lowers the learning curve for NetOps and DevOps staff, while also reducing the impact of, and rates for, human error.

WHY CHOOSE PLURIBUS NETWORKS' ADAPTIVE CLOUD FABRIC FOR YOUR DISTRIBUTED CLOUD?

Simply put, the Pluribus ACF is simple, adaptive, and open. Let's explore each of these attributes in more detail to show where and how these characteristics play out, and the value they can deliver.

- **The Pluribus ACF Is Simple.** All the switches that run the Netvisor ONE OS are federated and have a view of the entire state of the fabric, using tried-and-true distributed database technology for updates and concurrency controls. Thus, they function in the aggregate as a single logical fabric. This makes it simple to define and apply consistent configurations and policy controls across the entire fabric. Because each switch can see and control the entire fabric, there's no need to access a remote SDN controller for state information, telemetry data, to access end points, or to monitor and manage data flows.
- **The Pluribus ACF Is Adaptive.** Because the intelligence is distributed and control is available through any switch in the fabric, that fabric is incredibly adaptive. It works with equal facility in multiple network topologies, including spine and leaf in a multi-stage CLOS fabric topology, ring, and multi-site (or any combination of any or all of these types). In fact, the Pluribus ACF provides multi-site capability out of the box, with no need for adding (or even using) separate, complicated network controllers. It can also adapt to brownfield scenarios with existing legacy infrastructure because it uses standard networking protocols for the underlay.

> Simply put, the Pluribus ACF is simple, adaptive, and open.

- **The Pluribus ACF Is Open.** Companies and organizations that buy Netvisor ONE/ACF devices need not worry about vendor lock-in. The Pluribus ACF runs on whitebox hardware available from well-known vendors including Dell EMC, D-Link Systems, and Edgecore. Hardware and software are completely disaggregated, letting buyers swap out either or both over time, as market conditions, technical allegiances, and capabilities change.

CONTROLLERLESS NEXT-GENERATION SDN IS IDEAL FOR DISTRIBUTED CLOUD

One big advantage that comes from the Pluribus ACF is the absence of external SDN controllers. The Netvisor ONE OS and its ACF supports all capabilities of an SDN controller directly, but distributes the full state of the network and the SDN control intelligence to every switch node in the fabric.

This is important, because having no external controllers means reduced complexity and cost. Typically, each controller only supports n switches; and for each controller instance, industry best practices recommend three servers for redundancy. With those functions integrated into existing devices, making them an essential part of the fabric, there's no need to purchase servers or other specialized hardware on which to run controller software.

This also reduces the complexity that's inevitable if each edge site requires its own separate network controller, or if a hierarchy of controllers is required to tie core, distribution, and edge networks together.

In addition, external controllers typically use an OOB management channel, and a single management port on each switch can be a single point of failure. ACF uses in-band communication across a mesh of VXLAN tunnels, which means multiple paths and multiple ports with no single point of failure.

> Having no network controller greatly simplifies interconnecting and managing multiple sites.

Second, as edge sites are distributed, there can be serious latency penalties if nodes at a remote site need to communicate back to a centralized controller. This can defeat one of the key goals of lower latency championed by distributed cloud, and also can cause problems when rapid re-convergence is needed. The ACF re-converges much more quickly after link or node failures, because information-sharing and data propagation among switches in the fabric is much faster than with a centralized network controller.

Having no network controller greatly simplifies interconnecting and managing multiple sites, too. The ACF lets companies and organizations create and extend a distributed cloud across multiple geographical locations because there's no latency to reach back to a central controller—the controller intelligence is resident in every node in the fabric. Once the links are established and the nodes in the fabric are communicating with one another, you're done: that's all there is to setting up, managing, and maintaining a multi-site network.

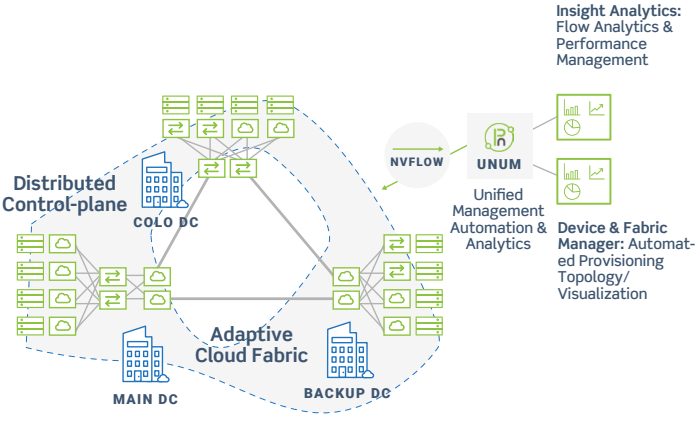

Figure 5: The Pluribus Adaptive Cloud Fabric builds one single logical network across multi-site data centers and edge locations with distributed intelligence, full visibility of all connected devices, and traffic to the port and flow level.

Because the fabric uses distributed database technology to hold network state and configuration information, controllers need pose no delays while learning new routes, propagating network change data or when responding to other devices' requests for such information.

PROGRAMMABILITY AND AUTOMATION

Because the ACF functions as a single federated fabric, programming applies to all switches equally (and quickly). Instead of having to remote into individual switches, programming changes can be made from any node via CLI or API, yet apply to all nodes.

The ACF supports a set of distributed, fabric-wide objects with an API that makes it easy to automate and innovate with network configuration, resource allocation, management, and more. Admins and programmers can work on one-off tasks or basic scripts through a powerful CLI (see the Netvisor ONE Technical Documentation, Command Reference 3.0.0 documents).

For heavier lifting, Netvisor ONE also supports a full-blown RESTful API (see the Netvisor ONE REST API 3.1.1 manual[1] for details on the self-documenting Swagger API tool; for the full set of API calls, see the complete REST API document[2]).

COMPREHENSIVE FABRIC-WIDE NETWORK SLICING/SEGMENTATION

Network slicing is mandated by the 3GPP, the premier standards body defining 5G architectural requirements. Netvisor ONE and ACF supports the industry's most comprehensive network slicing, also known as network segmentation. This environment creates full isolation across the management plane, control plane, and data

[1] https://techdocassets.pluribusnetworks.com/ONVL/3.1.1/Accton/PN-REST-API-311.pdf

[2] https://techdocassets.pluribusnetworks.com/ONVL/3.1.1/PN-Netvisor-REST-API-v311.pdf

plane. The network slice is fabric-wide and easily spans geographies. Enterprises or service providers can set up separate, non-intersecting virtual networks (VNETs) that can use overlapping VLANs and IP addresses to support multi-tenant scenarios.

Administrators can manage individual slices independently with their tool of choice, such as Ansible, CLI, Pluribus UNUM (Web-based GUI for management, automation, and analytics at the box level and fabric level) or any tool via the REST API.

A data center operator could, for instance, use this approach to isolate and insulate its IoT projects completely and separately from corporate traffic. The same goes for other highly distributed data and bandwidth-intensive applications such as video surveillance, for example.

Enterprises can also insulate and isolate separate lines of business from one another, or use the fabric to do likewise for classified and non-classified networks and associated data and computing resources.

The architecture of the Pluribus ACF supports easy definition, creation, and management for VNETs. The Netvisor ONE OS and ACF, in fact, permit admins to set up entirely non-overlapping networks—with separate and isolated routing tables, VLANs, IP addresses, resource domains, and more—so that multiple VNETs can happily and efficiently co-exist within a single switch fabric.

For each network slice, the same general programmability and ease of operation applies, just as it does to the fabric as a whole. Here again, there's no need for a distant node to communicate via an OOB channel back to an SDN controller to set up, configure, or maintain a slice. Any fabric node offers easy access to control, manage, and operate the slice across the entire fabric.

PERVASIVE FABRIC-WIDE VISIBILITY

Netvisor ONE and the ACF provide deep visibility into ports, flows, and connected devices on a per-slice basis. This rich telemetry is fabric-wide and can be pumped into a database that can hold up to 2 billion flows, which can then be forensically examined. Admins can dig into a slice to examine endpoint status and behavior. They can also

Enterprises or service providers can set up separate, non-intersecting virtual networks (VNETs) that can use overlapping VLANs and IP addresses to support multi-tenant scenarios.

examine traffic and data flows to monitor performance and reliability, run analytics, or speed troubleshooting and problem resolution practices and procedures.

Insight Analytics is a powerful graphical analytics tool that can be deployed in conjunction with ACF. This tool provides a visualization of all flow data over time and can be thought of as a network DVR, enabling network operators to go back in time to analyze, for example, network slowdowns over a particular time period and determine cause and effect.

PROVEN MISSION-CRITICAL DATA CENTER DEPLOYMENTS

Pluribus Networks Netvisor ONE OS and ACF are already widely deployed in some of the world's largest service providers, mobile carriers, and enterprises. In fact, through a partnership with Ericsson, Netvisor ONE OS and ACF have been deployed as the network fabric for Ericsson's Software Defined Infrastructure (SDI) and Network Function Virtualization Infrastructure (NFVi) in the virtualized mobile packet cores of many tier-1 mobile carriers. The platform has shown itself effective and robust in production use in these mission-critical environments.

MEET THE PLURIBUS PRODUCT/ SOLUTIONS PORTFOLIO

The Pluribus family of products supports the Adaptive Cloud Fabric and makes the distributed cloud a reality.

- **Netvisor ONE Network OS.** The Netvisor ONE Network OS is Pluribus Networks' primary and principal product. Based on The Linux Foundation's open source FRRouting software, Netvisor is an open, secure, and programmable

next-generation network OS. It's purpose-built to optimize the power and performance of bare-metal open networking hardware. Deployed in mission-critical enterprise and carrier networks, Netvisor offers powerful performance, complete network virtualization, and outstanding reliability and flexibility at cloud scale.

- **Adaptive Cloud Fabric (ACF).** This is enabled by the ability of Netvisor ONE to unify and aggregate all switches running Netvisor ONE into a single, virtual networking environment, with unified management, automation, visibility, and monitoring. The ACF is a holistic, distributed network architecture that uses a controllerless SDN approach for the management fabric and VXLAN tunnels for network virtualization, while using standard Layer 2 and Layer 3 protocols for the underlay. The ACF offers an adaptable and elastic networking environment that functions at cloud scale. Network services are fully integrated, and the ACF's ability to seamlessly interconnect multiple sites and locations supports massive scale-out, as well as scale-across capability. Read more at the Pluribus Adaptive Cloud Fabric overview page.[3]
- **Telemetry.** Pluribus Netvisor includes built-in telemetry capability, designed to feed rich, complete visibility on what's happening in your networks. Admins can see and use flow-based information for bare-metal and virtualized workloads across the entire network fabric, in addition to all endpoints connected to the fabric. Flows can be fed into a database that holds 2 billion flows for analysis and troubleshooting. Check out Pluribus Netvisor Fabric Visibility[4] for more information and examples.
- **UNUM.** A web-based tool that provides complete, GUI-based oversight for individual switches, as well as ACF and its components. Admins can use UNUM for management of fabric configurations and policies; automation of provisioning and deployment; and creation of network topology visualizations. UNUM also supports the Insight Analytics module, which provides performance management in real time, plus access to historical data for troubleshooting and capacity planning. See the UNUM datasheet[5] for more info, or click the UNUM search link here.[6]
- **Third-Party Partnerships.** Pluribus software is available pre-installed on a variety of whitebox and name-brand equipment through the company's partners. Key partners include Ericsson, Tibco, Dell EMC, D-Link Systems, Edgecore Networks, and more. See the Pluribus Technology Partners page for details. Pluribus also integrates with and interoperates with key orchestration solutions like vCenter from VMware, Open Daylight, and OpenStack.
- **Proven Solutions.** Pluribus platforms and tools are deployed at more than 200 companies and organizations around the globe, and at more than 50 mobile carriers, network operators, and service providers.

WHAT NEXT?

With so many enterprises and service providers extracting a competitive edge from Pluribus platforms and software, isn't it time your organization did the same? Pluribus can help you modernize your single-site or multi-site data center environment and simplify your deployment of distributed cloud. For more information on Pluribus, visit its website.[7] For more information, or for a sales contact with Pluribus or one of its partners, please reach out at pluribusnetworks.com/contact-us/.

[3] https://www.pluribusnetworks.com/assets/PluribusAdaptiveCloudFabric-Overview-032818.pdf

[4] https://www.pluribusnetworks.com/resources/pluribus-netvisor-fabric-visibility-2/

[5] https://www.pluribusnetworks.com/assets/Pluribus-UNUM-DS-053018.pdf

[6] https://www.pluribusnetworks.com/?s=UNUM

[7] https://www.pluribusnetworks.com/partners/technology-partners/