Technical Overview for:

BaaT

Blockchain-as-a-Transport

An Overview of BaaT – Blockchain-as-a-Transport

Described is **Blockchain-as-a-Transport (BaaT).**
BaaT is a software supporting network architecture that ensures efficient and secure data transport.
It is deployed as a service across certified hardware, but the technology is software-based and hardware agnostic. There are two main approaches to how BaaT is deployed:

Encrypted Data Pipeline
Any layer 2 or layer 3 traffic passing thru a BaaT Secure Data Pipeline is fully encrypted leveraging three distinct security methods:

A) Transport Layer Security via Blockchain **SHA-256** – Secure Hash Algorithm

B) Data Packet Encryption via **AES-256** – Advanced Encryption Standard

C) Post-Quantum Cryptography Standardization - (**PQCS**) realized via – BaaT Encryption's approach to QRNG, QKD and QRA.

The combination of the above three methods (SHA-256, AES-256, PQCS) ensures that neither the endpoints (source/destination) or the data payload can be penetrated or compromised. Encrypted Data Pipelines can be deployed in Point-to-point, Multi-Point, Hub and spoke, Fully meshed topologies.

NOTE: QRNG (Quantum Random Number Generation), QKD (Quantum Key Distribution) and QRA (Quantum Resistant Algorithm) are based on the latest standards being identified by NIST (National Institute of Standards and Technology). BaaT Encryption can TODAY support all the "first track" seven algorithm finalists announced on July 22, 2020 by NIST. This is important for governments and enterprises that want to be quantum compute secure. IBM has for example stated that "in five years" (2025), the effects of quantum computing will reach beyond the research lab.

Encrypted Network Overlay
The same layers of encryption existing for Encrypted Data Pipeline apply here. In this mode BaaT can connect geographically dispersed Layer 2 (L2) islands over any available infrastructure. Supportable infrastructures include leased lines, private fiber, Satellite, MPLS, IPVN, and the public Internet. It is the ability to support the public Internet that provides the greatest security and cost advantages (see Hacking and MITM Concerns).

BaaT solves many of the disadvantages of other network overlay technologies by integrating blockchain with VXLAN. VXLAN was originally created without a control plane, and as such the scope is normally limited to a single data center or cloud.

BaaT greatly enhances the operation of VXLAN by adding a control plane component to it and extending the VXLAN working domain beyond the boundary of a local data center or even a public cloud.

The public Internet is considered unstable for protocols that require deterministic data paths – BaaT overcomes this challenge by leveraging blockchain as the control plane element. Source and destination endpoints learn from the distributed ledger how to put packets together efficiently in support of all kinds of network traffic including: Unicast, Multicast, and Broadcast.

BaaT Use Case Scenarios

Project 1: Secure Video Conferencing
Mode: **Encrypted Data Pipeline**
Network: IPVPN 10mbs leased circuits
Design: Hub/Spoke – 2 locations

DATE:	**Tuesday, December 8th, 2020**		
Start:	**9:00**	End:	**11:00**
Project Manager:	Mayande Walker		
Location 1:	Kabanbai Bayr 28 (Nur-Sultan)	Location 2:	Kunaev 10 (Nur-Sultan)
Engineer:	Alexander Sinchenko	Resource:	Alexander Sinchenko
1 SIDE Tech Overview (SPOKE):		**2 SIDE Tech Overview (HUB):**	
SPOKE: 2x BaaT Blackboxes connected to 2x Transtelecom MPLS circuits with 10 mbps speed each.		HUB: 2x BaaT Blackboxes connected to 2x Transtelecom MPLS circuits with 10 mbps speed each.	
AVG BaaT SPEED:	**9.3 Mbps (on 10 Mbps MPLS circuit)**		
HUB<<>>SPOKE RESULTS:	2x BaaT PVC Circuits: UP. End-to-End Connectivity confirmed via PING and SSH tests: **SUCCESSFUL**. Encryption: **CONFIRMED** - File Transfer Tests: **SUCCESSFUL**		
SPOKE<<>>SPOKE RESULTS:	1x BaaT SVC Circuit: **UP**. End-to-End Connectivity confirmed via PING and SSH tests: **SUCCESSFUL**. Encryption: **CONFIRMED** - File Transfer Tests: **SUCCESSFUL - VIDEO CONFERENCE TEST: SUCCESSFUL**		

Details: Carrier provided 2x 10 Mbps MPLS private circuits across two office locations. 4x BaaT Edge Black boxes were deployed in a Hub/Spoke model. Three (3) BaaT overlay circuits were built and deployed – 2x BaaT PVC across the MPLS and 1x BaaT SVC to logically connect the two spokes within the design. Testing focused on "full mesh" capabilities between the two spokes – Spoke 1 and Spoke 2 which by virtue of design traverses 2x BaaT PVC circuits – Hub 1 to the West and Hub 2 to the East as shown in the below diagram.

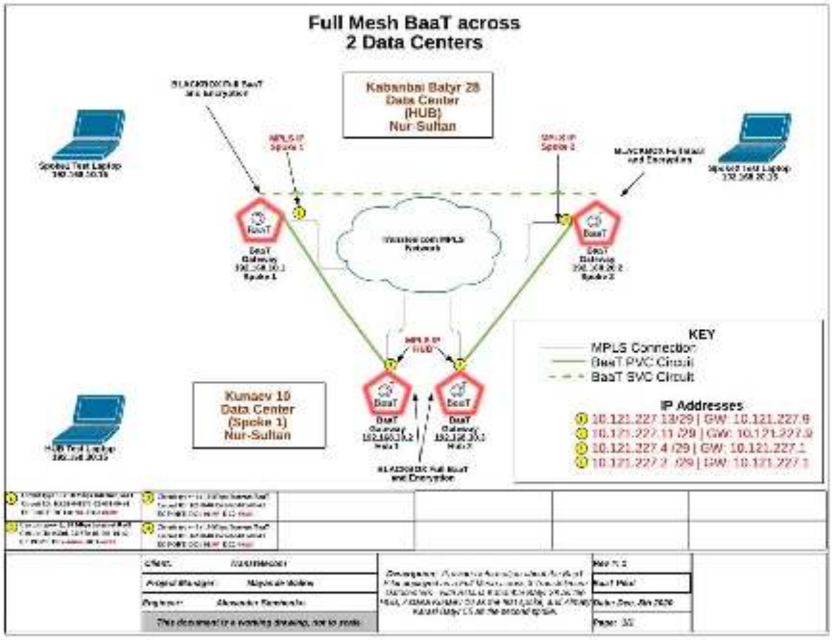

Full Mesh BaaT across
2 Data Centers
Kabanbai Batyr 28
Data Center
(HUB)
Nur-Sultan
Kunaev 10
Data Center
(Spoke 1)
Nur-Sultan
BaaT
BaaT
BaaT
BaaT
KEY
MPLS Connection
BaaT PVC Circuit
BaaT SVC Circuit
IP Addresses
① 10.121.227.13/29 | GW: 10.121.227.9
② 10.121.227.11 /29 | GW: 10.121.227.9
③ 10.121.227.4 /29 | GW: 10.121.227.1
④ 10.121.227.2 /29 | GW: 10.121.227.1
This document is a working drawing, not to scale

Data Flow: (see Diagram) – BaaT Encrypted Data Pipeline was deployed as a Mesh across 2 locations (See BaaT Encryption Key Management and Encrypted BaaT Encapsulation Flow). A BaaT Edge "Black Box" served as the CPE (Customer Premise Equipment) with IPVPN Circuits connected to the WAN port and testing laptops connected to the LAN port. A client server video conferencing application (TruConf) was used on the laptops at each location to test performance of the encrypted channels. TCPDump and Wireshark were used to show that all packets traversing these channels were fully encrypted.

All packets passing thru WAN ports are fully encrypted and de-encrypted within the BaaT Edge device as it traverses the LAN port. Blockchain ledger provides BaaT Node to Node security as well as the instruction set to the Linux kernel of the BaaT Edge device regarding authenticated networks and subnets.

Certification Testing:
Encryption: Should include capturing packets traversing the logical pipeline between WAN interfaces of the BaaT Edge device. TCPDump and Wireshark can provide packet level verification of encryption. More intrusive testing could be done with a physical packet sniffer or switch that can support port mirroring in order to capture the data flow for offline analysis of the cryptography algorithm strength (currently set at AES-256 and including PQCS-QRA).

Project 2: BaaT SAT (Satellite Circuit)
Mode: **Encrypted Network Overlay**
Network: 2mbps CIR/MIR GEO (KazSat 2) via IDirect & 20 Mbps Internet Connection
Design: Point-to-Point – 2 locations

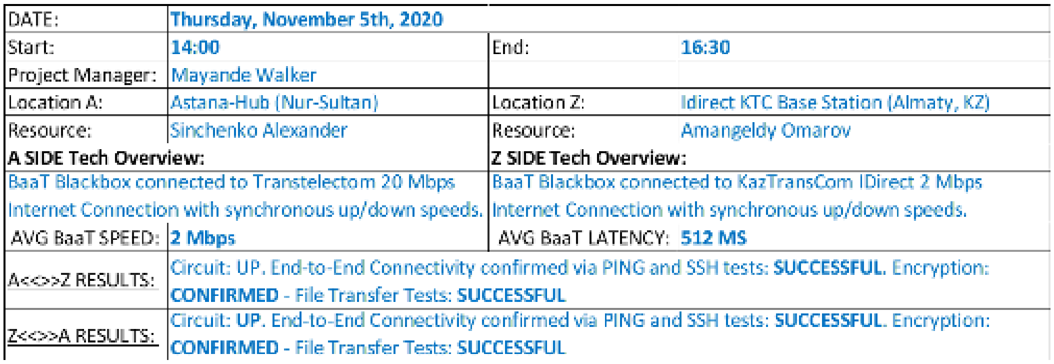

DATE:	**Thursday, November 5th, 2020**		
Start:	**14:00**	End:	**16:30**
Project Manager:	Mayande Walker		
Location A:	Astana-Hub (Nur-Sultan)	Location Z:	Idirect KTC Base Station (Almaty, KZ)
Resource:	Sinchenko Alexander	Resource:	Amangeldy Omarov
A SIDE Tech Overview:		**Z SIDE Tech Overview:**	
BaaT Blackbox connected to Transtelectom 20 Mbps Internet Connection with synchronous up/down speeds.		BaaT Blackbox connected to KazTransCom IDirect 2 Mbps Internet Connection with synchronous up/down speeds.	
AVG BaaT SPEED:	**2 Mbps**	AVG BaaT LATENCY:	**512 MS**
A<<>>Z RESULTS:	Circuit: UP. End-to-End Connectivity confirmed via PING and SSH tests: **SUCCESSFUL.** Encryption: **CONFIRMED** - File Transfer Tests: **SUCCESSFUL**		
Z<<>>A RESULTS:	Circuit: **UP.** End-to-End Connectivity confirmed via PING and SSH tests: **SUCCESSFUL.** Encryption: **CONFIRMED** - File Transfer Tests: **SUCCESSFUL**		

Details: BaaT Encrypted Network Overlay was deployed across 2 locations. On one side was a Satellite based connection providing 2 mbps synchronous upload/download Internet towards a second connection with 20 mbps synchronous Internet. This required no custom fiber or private circuits between base station and client data center. Cost effective and with 2 hours provisioning time.

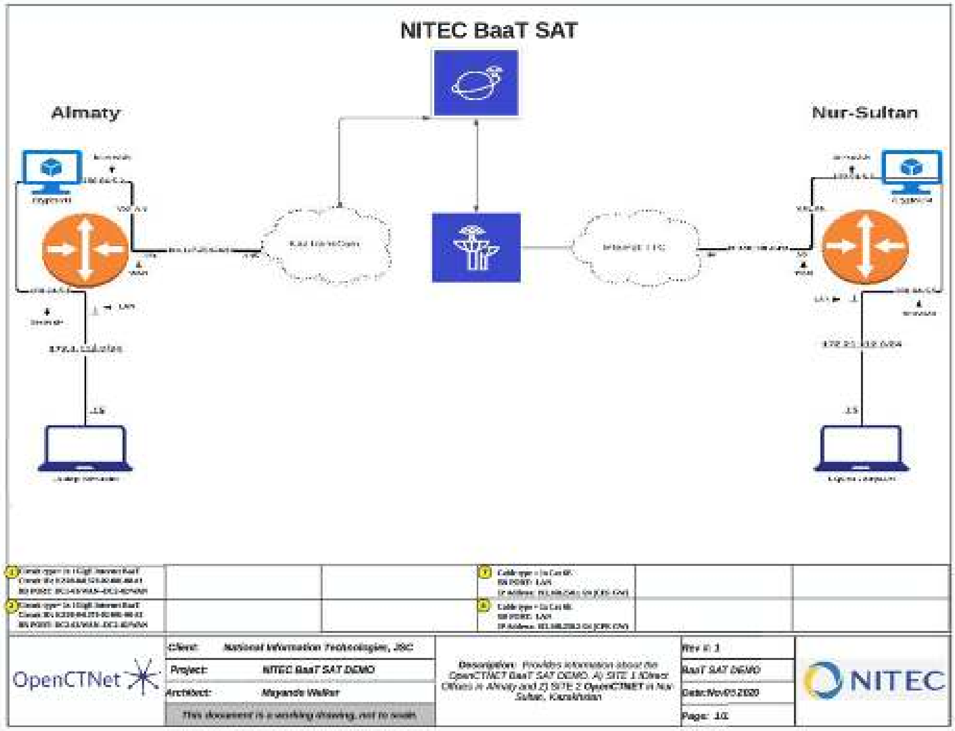

NITEC BaaT SAT
Almaty
Nur-Sultan
VXLAN
WAN
LAN
.15
OpenCTNet
Client: National Information Technologies, JSC
Project: NITEC BaaT SAT DEMO
Architect:
This document is a working drawing, not to scale.
Description: Provides information about the OpenCTNET BaaT SAT DEMO. A) SITE 1 Offices in Almaty and 2) SITE 2 OpenCTNET in Nur-Sultan, Kazakhstan
Rev #: 1
BaaT SAT DEMO
Date:Nov05 2020
Page: 1/1
NITEC

<u>Data Flow:</u> (see Diagram) – To create the virtual circuit VXLAN protocol was engaged along with BaaT tunneling technology was used to extend the VXLAN 1 hop limitation – (See VXLAN Overview, BaaT VXLAN Control Plane Operation, and Encrypted BaaT Encapsulation Flow). All packets passing thru WAN ports are fully encrypted and de-encrypted within the BaaT Edge device as it traverses the LAN port. Blockchain ledger provides BaaT Node to Node security as well as the instruction set to the Linux kernel of the BaaT Edge device regarding authenticated networks and subnets.

<u>Certification Testing:</u>

Circuit: Should include stability and performance of data file transfers, and various applications across the secured tunnels and up to the bandwidth specified (2 mbps).

Security: Verification that the internal networks are hidden/encrypted. Verification that the external networks are "whitelisted" - that the ISP/Carrier has limited communication to the external IP address to only approved networks.

Encryption: Should include capturing packets traversing the logical pipeline between WAN interfaces of the BaaT Edge device. TCPDump and Wireshark can provide packet level verification of encryption. More intrusive testing could be done with a physical packet sniffer or switch that can support port mirroring in order to capture the data flow for offline analysis of the cryptography algorithm strength (currently set at AES-256 and including PQCS-QRA).

ETS "GO": That based on the security and encryption of the circuit that it can be certified for ETS GO networks.

Project 3: DCI Data Transport
Mode: **Encrypted Network Overlay**
Network: 6x 1 Gb DCI Channels support by 12x BaaT Circuits with diverse ISP for failover and redundancy
Design: Hub/Spoke – 4 sites Full Mesh capability (Nur-Sultan – Hub; Almaty, Karaganda, Petropavlovsk.

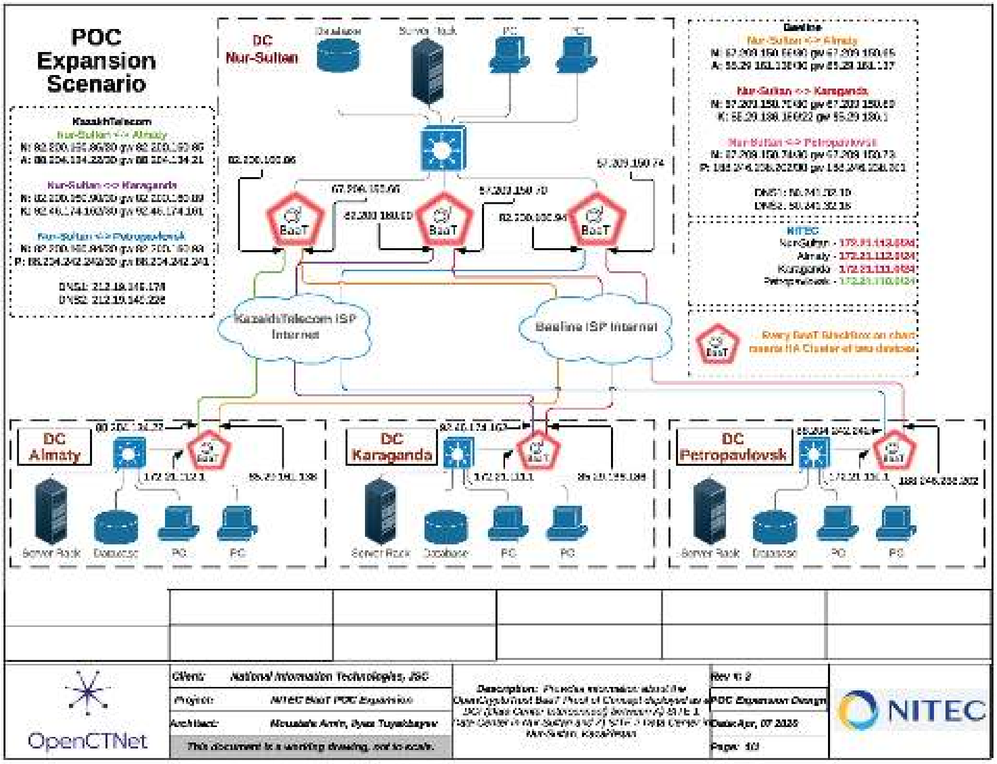

Data Flow: (see Diagram) – To create the virtual circuits VXLAN protocol was engaged along with BaaT tunneling technology was used to extend the VXLAN 1 hop limitation – (See VXLAN Overview, BaaT VXLAN Control Plane Operation, and Encrypted BaaT Encapsulation Flow). All packets passing thru WAN ports are fully encrypted and de-encrypted within the BaaT Edge device as it traverses the LAN port. Blockchain ledger provides BaaT Node to Node security as well as the instruction set to the Linux kernel of the BaaT Edge device regarding authenticated networks and subnets.

Certification Testing:
Circuit: Should include stability and performance of data file transfers, and various applications across the secured tunnels and up to the bandwidth specified (500 mbps).
Security: Verification that the internal networks are hidden/encrypted. Verification that the external networks are "whitelisted" - that the ISP/Carrier has limited communication to the external IP address to only approved networks.
Encryption: Should include capturing packets traversing the logical pipeline between WAN interfaces of the BaaT Edge device. TCPDump and Wireshark can provide packet level verification of encryption. More intrusive testing could be done with a physical packet sniffer or switch that can support port mirroring in order to capture the data flow for offline analysis of the cryptography algorithm strength (currently set at AES-256 and including PQCS-QRA).
ETS "GO": That based on the security and encryption of the circuit that it can be certified for ETS GO networks.

Encrypted BaaT Encapsulation Flow

In the case of Encrypted Network Overlay it is helpful to have a clear understanding of the Blockchain Control Plane Abstraction process (see below slide). In general, there are 3 phases in which a layer 2 frame becomes VXLAN capable, layer 3 capable, and then abstracted into two separate encrypted packets. One packet is destined for the blockchain ledger for control plane information distribution. The second is sent thru the Internet via UDP and is unable to be re-assembled without the control plane information that is learned by the destination node via the blockchain ledger.

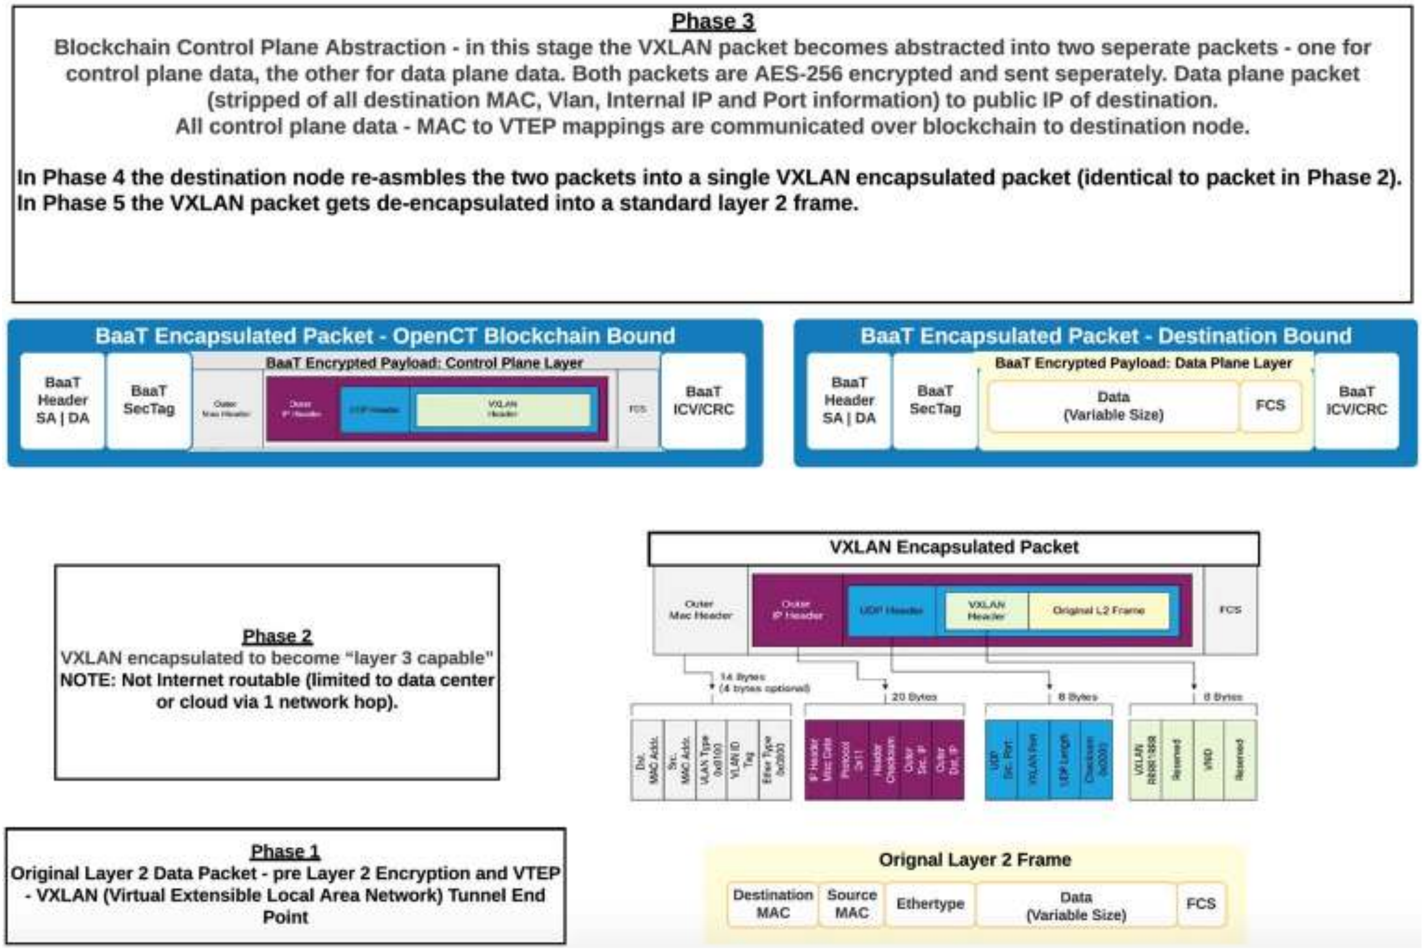

Phase 3

Blockchain Control Plane Abstraction - in this stage the VXLAN packet becomes abstracted into two seperate packets - one for control plane data, the other for data plane data. Both packets are AES-256 encrypted and sent seperately. Data plane packet (stripped of all destination MAC, Vlan, Internal IP and Port information) to public IP of destination. All control plane data - MAC to VTEP mappings are communicated over blockchain to destination node.

In Phase 4 the destination node re-asmbles the two packets into a single VXLAN encapsulated packet (identical to packet in Phase 2).
In Phase 5 the VXLAN packet gets de-encapsulated into a standard layer 2 frame.

Phase 2

VXLAN encapsulated to become "layer 3 capable"
NOTE: Not Internet routable (limited to data center or cloud via 1 network hop).

Phase 1

Original Layer 2 Data Packet - pre Layer 2 Encryption and VTEP - VXLAN (Virtual Extensible Local Area Network) Tunnel End Point

BaaT Encryption Key Management

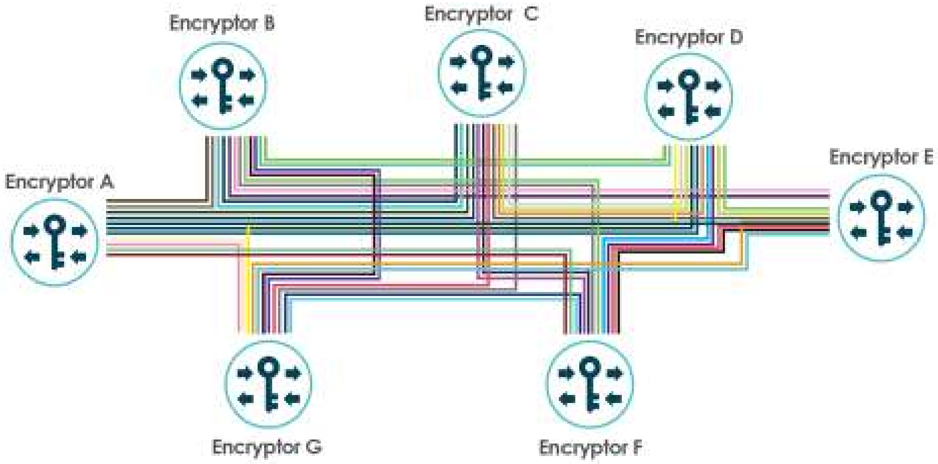

Secure Keys

Secure encryption depends on key security in all phases:

- A safe key requires entropy (randomness). Whilst both hardware and software can be used to create random numbers, true randomness comes from a hardware source.
- The encryption keys are in plain text while encrypting and decrypting are dependent on a safe encryption environment. BaaT encryptors feature a tamper-proof enclosure. Any tampering inevitably leads to the zeroisation of all data in memory, including the keys in use.
- Keys need to be secure while being transported between encryptors. Keys are always encrypted while in transit.
- Each encryptor has its own certificate, issued by the Certificate Authority (CA). The certificate is, initially, secret. Its public key is used as a digital signature, so the recipient can verify the sender.
- The key exchange uses the certificate to sign the keys (or partial keys) that are exchanged, ensuring that they are coming from the correct remote device.
- The partial keys are generated completely inside the encryptor, without any user having access to it. After exchanging the partial keys, both sides calculate the same shared secret.
- Subsequently, the encryptor internally generates the master key and encrypts it with the shared secret. The encryptor also generates the session key and uses the master key to encrypt it. The transmission of the master and session keys from one encryptor to the other is always encrypted.

Multipoint-to-Multipoint Key Management System

- An encrypted high-speed data network's robustness is in a large part determined by the type of encryption key management. Ensuring that the transmitted data's encryption keys are not accessible to an unauthorized party is essential to ensuring that a network breach **will only result in meaningless data in the hands of an unauthorized party**.
- In this scenario, BaaT key management supports the combination of pairwise key system with a distributed group key system. This allows multicast frames to be encrypted with a group key when using a point-to-multipoint topology. For unicast frame encryption, pairwise keys are used; while multicast frames are encrypted using the group key system. As multicast groups can originate at different sites, the encryptor at such sites must be able to act as group key server for the multicast groups originating there.

- For unicast frame encryption each encryptor has a table with the local and remote MAC addresses. Security associations are created automatically. The MAC addresses are automatically learned and assigned to the security associations using a proprietary auto-discovery protocol. Multicast groups can either be discovered in discovery mode and be added automatically, or they can be added manually. The key servers work independently of each other and thus offer redundancy. Each multicast group has a senior member that is responsible for the key generation, the key exchange and the key updates. If this senior member drops out, the next member of the multicast group is promoted to senior member and takes over. The group key system ensures that the right keys are available at the right time to the group members. The group key server for that group generates the keys for that group, using a hardware-based true random number generator for the seed.

Encrypted Network Overlay – Blockchain Based Data Storage and Retrieval

Connecting L2 islands is not new. There are several popular solutions that fall under the "Overlay Networking Technologies" umbrella such as:

- **Virtual Extensible LAN (VXLAN)**
- Network Virtualization using Generic Routing Encapsulation (NVGRE)
- Overlay Transport Virtualization (OTV)
- Virtual Private LAN Service (VPLS)
- IEEE 802.1ah Provider Backbone Bridge (PBB)

Each involves encapsulating L2 frames within other headers either at L2 or L3, and each comes with its own pros and cons such as distance limitations, scaling problems, and management complications. BaaT solves many of the disadvantages of other overlays by integrating blockchain with VXLAN. VXLAN was originally drafted as an overlay technology that can work without a control plane. It has proven to be an overlay of choice, but its scope is normally limited to a single data center or cloud.

BaaT greatly enhances the operation of VXLAN by adding a control plane component to it and extending the VXLAN working domain beyond the boundary of a local data center or even a public cloud.

BaaT operation across the public Internet is appealing as a viable WAN option for many network operators such as enterprises, service providers, and telcos in front of conventional, expensive WAN options such as dedicated links, MPLS, or Virtual Private Networks (VPNs).

BaaT is useful for any critical high-frequency trading application as described above. These applications require many events and transactions to be recorded over the Blockchain while at the same time ensuring maximum stability, scalability, security, and requiring the fastest convergence time.

BaaT achieves control plane operation via Blockchain. See Figure 1.

In this mode, the VXLAN Tunnel Endpoints (VTEPs) are also nodes of a public or private Blockchain that can span the public Internet.

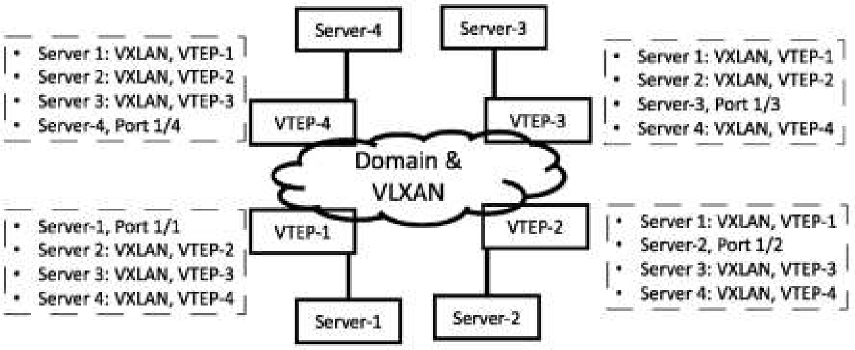

Figure 1.

The local MAC learning technique is the same as with any other VXLAN operation: The VTEPs learn the local MAC addresses via their local ports, as shown in Figure 2: BaaT Initial State, and then the addresses are advertised/published as reachable through their VTEP IPs over the Blockchain using transactions that are packed into proper blocks.

Steps to publish a stream of hexadecimal data over the Blockchain:

i. VTEP Converts Alphanumeric Text to Hexadecimal Text
ii. VTEP publishes the Hexadecimal Text over the Blockchain
iii. The other recipients VTEPs retrieve the Hexadecimal Text from over the Blockchain and convert it back to Alphanumeric Text
iv. The recipient VTEP uses this data for further communications with all other VTEPs

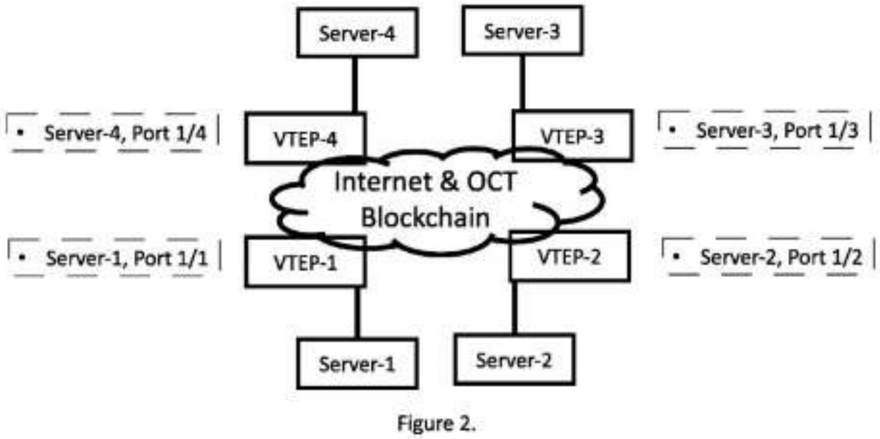

Figure 2.

Example:
Customer #1555 LAN segment is connected to VTEP-1 (in Figure 2), VTEP-1 needs to participate in VXLAN with VNID 10123. VTEP-1 just learned the Media Access Control (MAC) address (00-14-22-01-23-45) from one of the locally attached servers belonging to Customer#1555 (Server-1 attached to Port 1/1).

The VTEP-1 IP address is 10.1.1.178, and this is the IP address that other VTEPs need to use to reach VTEP-1. The message that will be published from that VTEP over the Blockchain is typically a MAC-to-VTEP mapping message that also includes the Customer ID as well as the VNID.

i. VTEP Converts Alphanumeric Text to Hexadecimal Text

1. From Alphanumeric Text: ‘customer 1555 vnid 10123 mac address 00-14-22-01-23-45 VTEP 10.1.1.178’
2. To Hexadecimal Text: ‘637573746f6d6572203135353520766e6964203130313233206d616320616464726573732030302d31342d32322d30312d32332d343520565445502031302e312e312e313738’

ii. VTEP Publishes Hexadecimal Text to Blockchain

iii. VTEPs see the Blockchain Hexadecimal Text and reads them back to all VTEPs

iv. VTEPs Converts Hexadecimal Text to Alphanumeric Text

1. From Hexadecimal Text: ‘637573746f6d6572203135353520766e6964203130313233206d616320616464726573732030302d31342d32322d30312d32332d343520565445502031302e312e312e313738’
2. To Alphanumeric Text: ‘customer 1555 vnid 10123 mac address 00-14-22-01-23-45 VTEP 10.1.1.178’

This message can be seen by all VTEPs participating in the Blockchain but only those VTEPs that are interested in Customer ID 1555 and VXLAN VNID 10123 will use this message, translate it, and add its content to their local copy of the MAC-to-VTEP mappings. Because the different MAC-to-VTEP mappings are distributed over the Blockchain to all participating nodes/VTEPs, as the final state shown in Figure 3.

- No data-plane learning is required for unknown unicast MAC addresses.
- No IP multicast underlay is required. This is why BaaT can span beyond the boundary of a data center or cloud to the public Internet.
- Because of the distributed nature of Blockchain, no significant delay is expected between the different nodes.
- For the broadcast and multicast traffic, the head-end replication is always the solution as in other control-plane-based VXLAN modes.

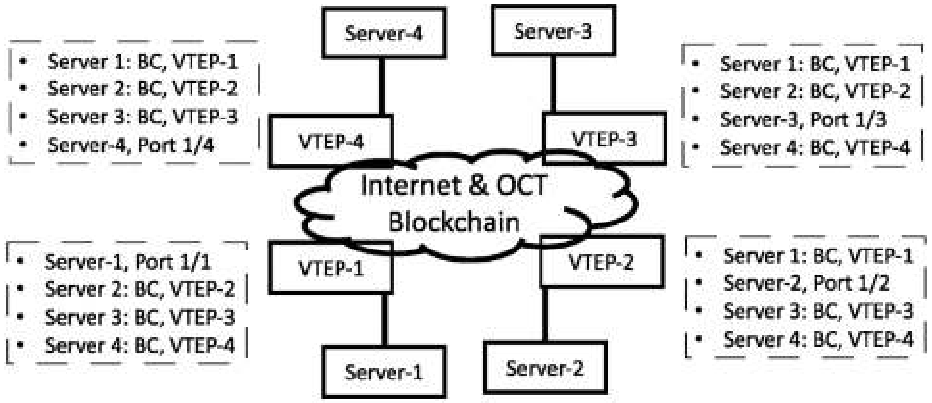

Figure 3.

VXLAN Overview

As its name indicates, Virtual eXtensible Local Area Network (VXLAN) is designed to provide the same Ethernet Layer 2 network services as Virtual LAN (VLAN) does today, but with greater extensibility and flexibility. Compared to VLAN, VXLAN offers the following benefits:

- *Flexible placement of multitenant segments throughout the data center.* VXLAN extends Layer 2 segments over the underlying shared network infrastructure so that tenant workloads can be placed across physical pods in the data center.
- *Higher scalability to address more Layer 2 segments.* VLANs use a 12-bit VLAN ID to address Layer 2 segments, which results in limited scalability of only 4094 VLANs. VXLAN uses a 24-bit segment ID known as the VXLAN Network IDentifier (VNID), which enables up to 16 million VXLAN segments to coexist in the same administrative domain.
- *Better utilization of available network paths in the underlying infrastructure*. VLAN uses the Spanning Tree Protocol (STP) for loop prevention, which wastes half of the network links by blocking redundant paths. In contrast, VXLAN packets are transferred through the underlying network based on its Layer 3 header and can take complete advantage of Layer 3 routing, equal-cost multipath (ECMP) routing, and link aggregation protocols to use all available paths.

VXLAN Format, Traffic Flow, and ECMP

VXLAN encapsulation adds 50 bytes to the original L2 frame by adding four headers:

- VXLAN header
- UDP header
- Outer IP header
- Outer MAC header

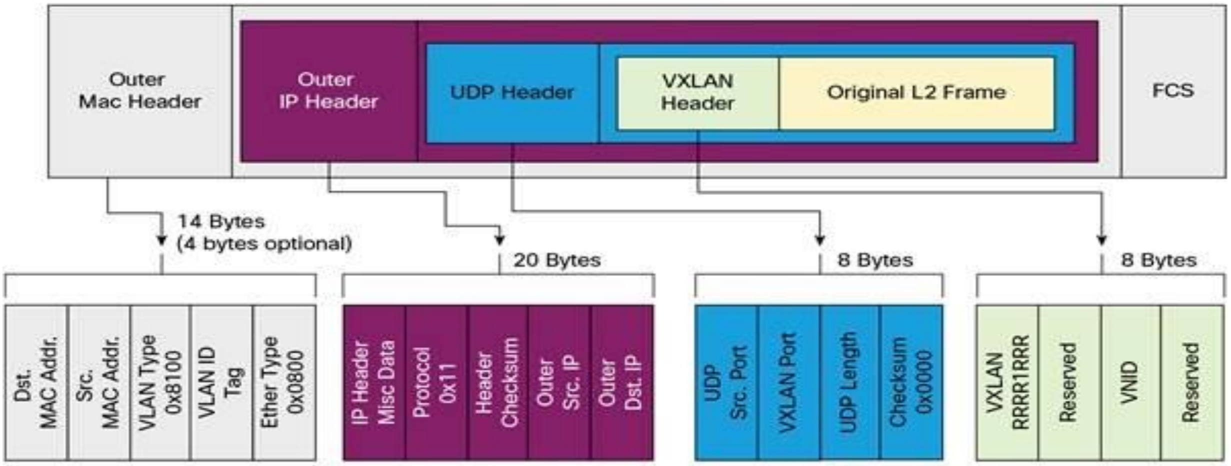

Figure 1.1: VXLAN Encapsulation

Encapsulated VXLAN packets are forwarded between VTEPs (VXLAN Tunnel End Points) based on the native forwarding decisions of the transport network. VXLAN was originally created to be bounded inside the data center where the underlay transport networks are designed and deployed with multiple redundant paths and take advantage of various multipath load-sharing technologies to distribute traffic loads on all available paths. It is desirable to share the load of the VXLAN traffic in the same fashion in the transport

network.
A typical VXLAN transport network is a routed (L3) IP network that uses the standard IP Equal-Cost Multipath (ECMP) to balance the traffic load among multiple best paths. To avoid out-of-sequence packet forwarding, flow-based ECMP is commonly deployed. An ECMP flow is defined by the source and destination IP addresses and optionally the source and destination TCP or UDP ports in the IP packet header. Because all of the VXLAN packet flows between a pair of VTEPs have the same outer source and destination IP addresses, and all VTEP devices must use one identical destination UDP port that can be either the Internet Allocated Numbers Authority (IANA)allocated UDP port 4789 (or a customer-configured port), the only variable element in the ECMP flow definition that can differentiate VXLAN flows from the transport network standpoint is the source UDP port. A similar situation for Link Aggregation Control Protocol (LACP) hashing occurs if the resolved egress interface based on the routing and ECMP decision is a LACP port channel. The LACP uses the outer VXLAN packet header for link load-share hashing, which results in the source UDP port being the only element that can uniquely identify a VXLAN flow.

VXLAN Modes of Operation

From the various industry implementations of VXLAN, we can categorize the VXLAN modes of operation into two main categories, each with two sub-categories:

- Control-Plane-LessVXLAN
  - Control-Plane-LessMulticastVXLAN
  - Control-Plane-LessUnicastVXLAN
- Control-PlaneVXLAN
  - Controller-Based VXLAN
  - EVPN VXLAN

The differences are mainly in the underlay transport network multicast capability, how to deal with Broadcast, Unknown Unicast & Multicast (BUM) traffic as well as the method of discovery and distribution of MAC addresses.

Control-Plane-Less Multicast VXLAN

As its name implies, there is no control or signaling established prior to the VXLAN operation.
This mode is according to the original VXLAN specification in RFC7348.
This mode requires the underlay transport network to fully support IP multicast,and every VTEP node to join the proper multicast domain.

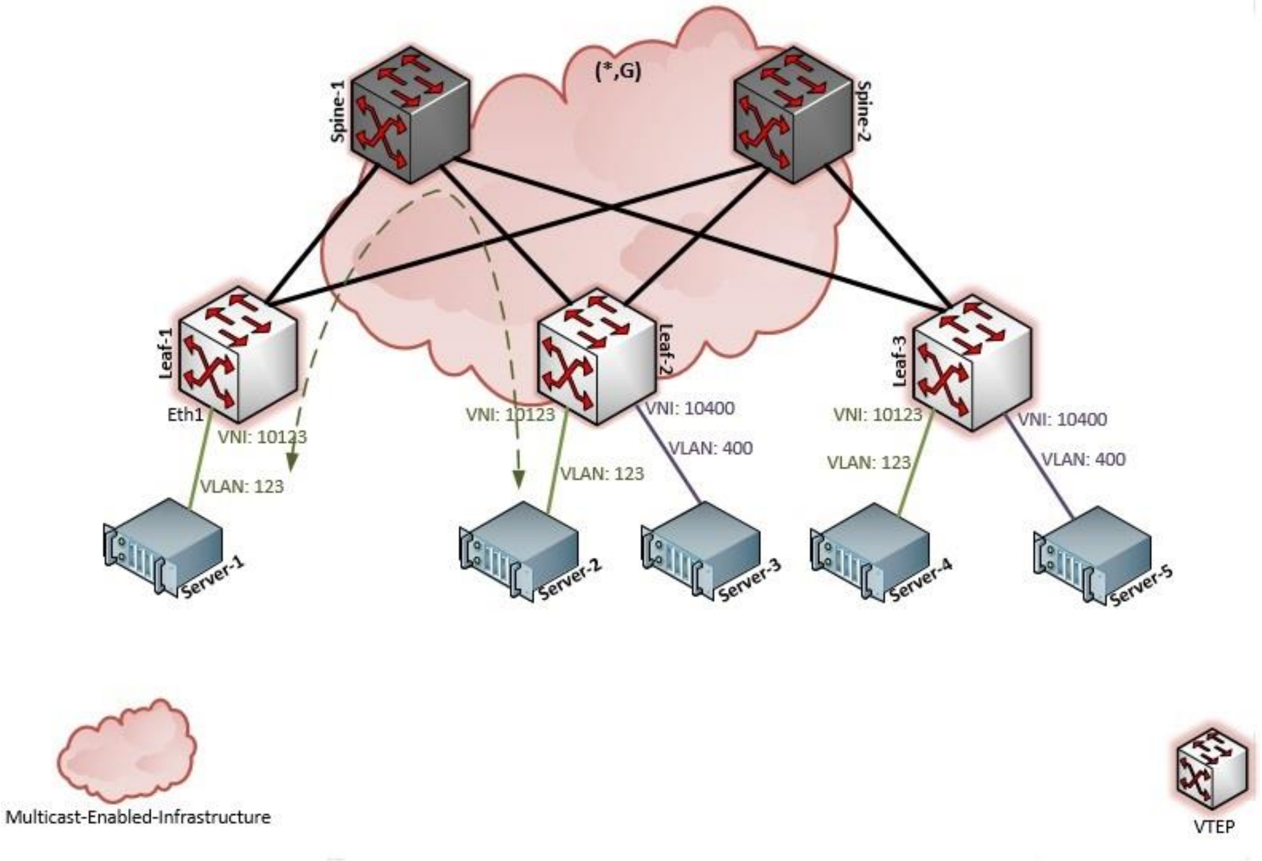

Figure 2.1: Control-Plane-Less Multicast VXLAN

In this mode, the BUM traffic is always carried over multicast. Data plane (or flow-based) learning is based on the "flood and learn" technique in which the remote VTEPs learn a MAC address because of the conversational MAC address learning approach:

- The destination VTEP learns the inner source MAC of any received VXLAN IP packet (for example a broadcasted ARP request message carried over multicast).
- The source MAC address is then mapped to the source VTEP that originated the VXLAN packet.
- The Originating VTEP learns the remote MAC address to VTEP mapping once it receives the VXLAN encapsulated unicast ARP reply message from the receiving VTEP.
- All subsequent traffic to a known MAC address will be unicast IP encapsulated VXLAN.

Control-Plane-Less Unicast VXLAN

Like the Control-Plane-Less-Multicast VXLAN, there is no control or signaling established prior to the VXLAN operation. Instead, a list of all available and participating VTEPs are configured on each VTEP per supported VXLAN.

In this mode, the underlay transport network doesn't need to support IP multicast.

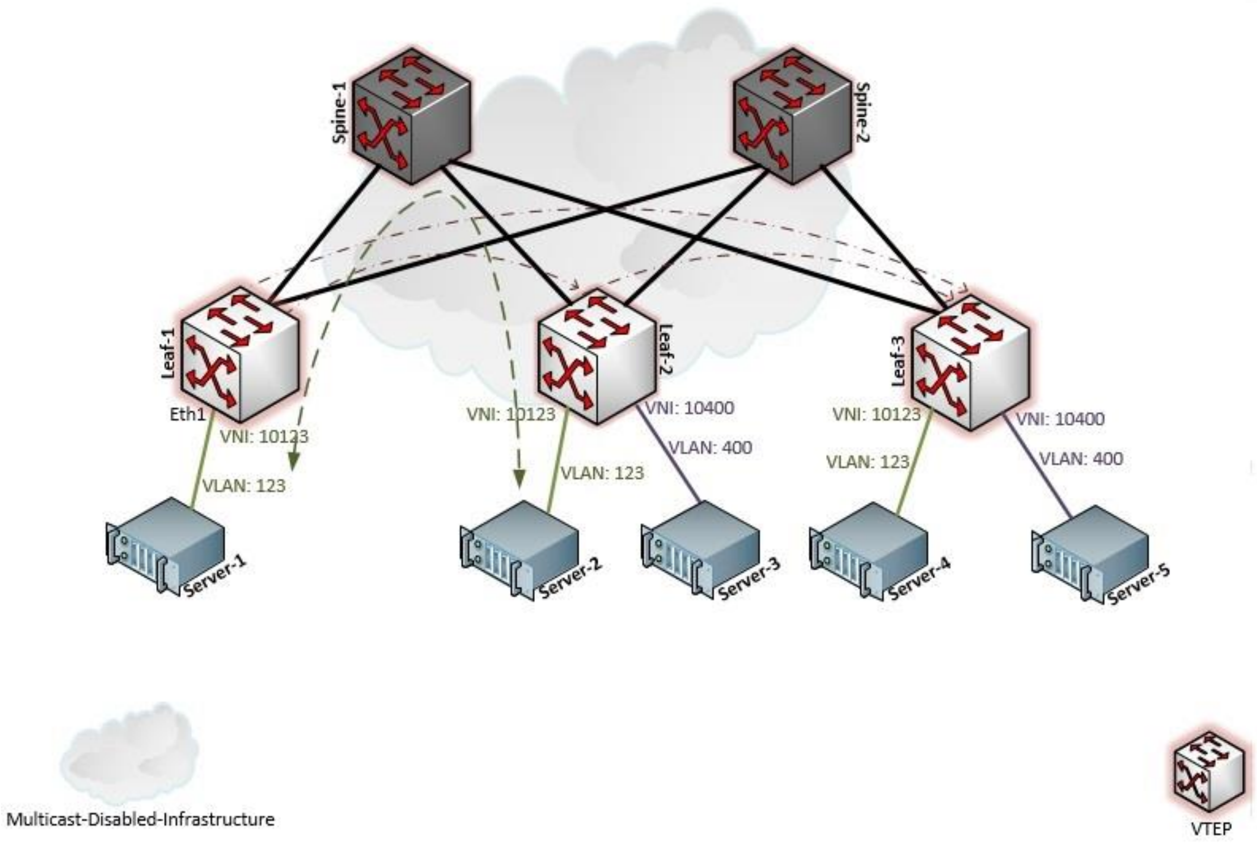

Figure 3: Control-Plane-Less Unicast VXLAN

Instead multicasting BUM traffic through the underlay transport network as in the previous mode of operation, head-end replication is used. The originating VTEP replicates the VXLAN packet and sends a copy to every other VTEP participating in this same VXLAN.

The list of VTEPs must be configured, changed, and updated manually on every VTEP in the VXLAN domain.

The data plane learning technique as described in the previous section is also used in this mode of operation.

Control-Plane VXLAN

In control-plane modes, there is no need for IP multicast in the underlay transport network. Head-end replication ise used, as in the previous mode, for broadcast and multicast traffic that is to be sent to all VTEPs.

Dealing with unknown unicast traffic is what differentiates this mode of operation from the previous modes. In this mode, a control plane exists to distribute the MAC-to-VTEP mapping entries between the different VTEPs; hence there is no need for any data plane learning technique.

This control plane piece could be a Controller such as VMware NSX, Midokura, Nuage, and Openstack; a signaling protocol such as MP-BGP inEVPN-based VXLAN; or a blockchain as in our proposed BaaT.

Controller-Based VXLAN

Data-Plane learning is optional or even not needed in controller-based VXLAN. The controller synchronizes all the MAC addresses as soon as the different switches learn them from their local ports.

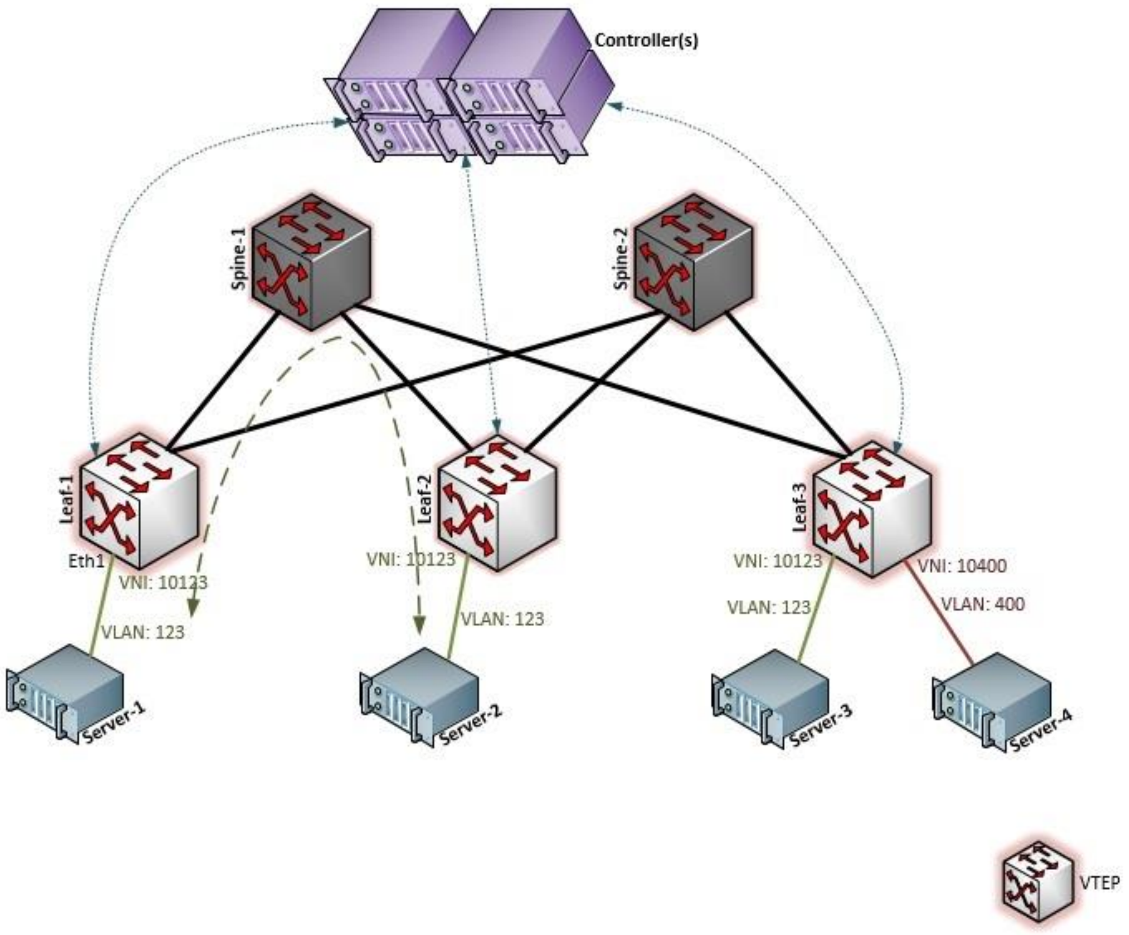

Figure 4: Controller-Based VXLAN

In Figure5, Leaf-1 learns the MAC address of Server-1 from its local port Eth1. This information is automatically and immediately synchronized with the controller that in turn pushes that info to Leaf-2, Leaf-3,and any other VTEP in the same VXLAN domain.

This VXLAN operation depends on the distribution of all learned MAC addresses from the different VTEPs via the controller that pushes to all VTEPs a complete -- and always updated-- list of MAC-to-VTEP mapping entries.

Because of that, there is no unknown unicast. The list of all communicating MACs is always present and updated on each VTEP. In the case of an unknown MAC – such as a destination outside the local VXLAN domain -- the local VTEP can direct it via the default entry towards the VXLAN gateway.

For broadcast and multicast traffic, head-end replication is always the solution.

EVPN VXLAN

In EVPN-VXLAN, each VTEP is a Provider Edge (PE) node and learns the local MAC addresses associated to its VXLANs from its local ports.Using the Multi-Protocol BGP (MP-BGP) EVPN address family, these entries are propagated to all other PEs either via direct MP-BGP sessions or through BGP route reflectors as shown in Figure 6.

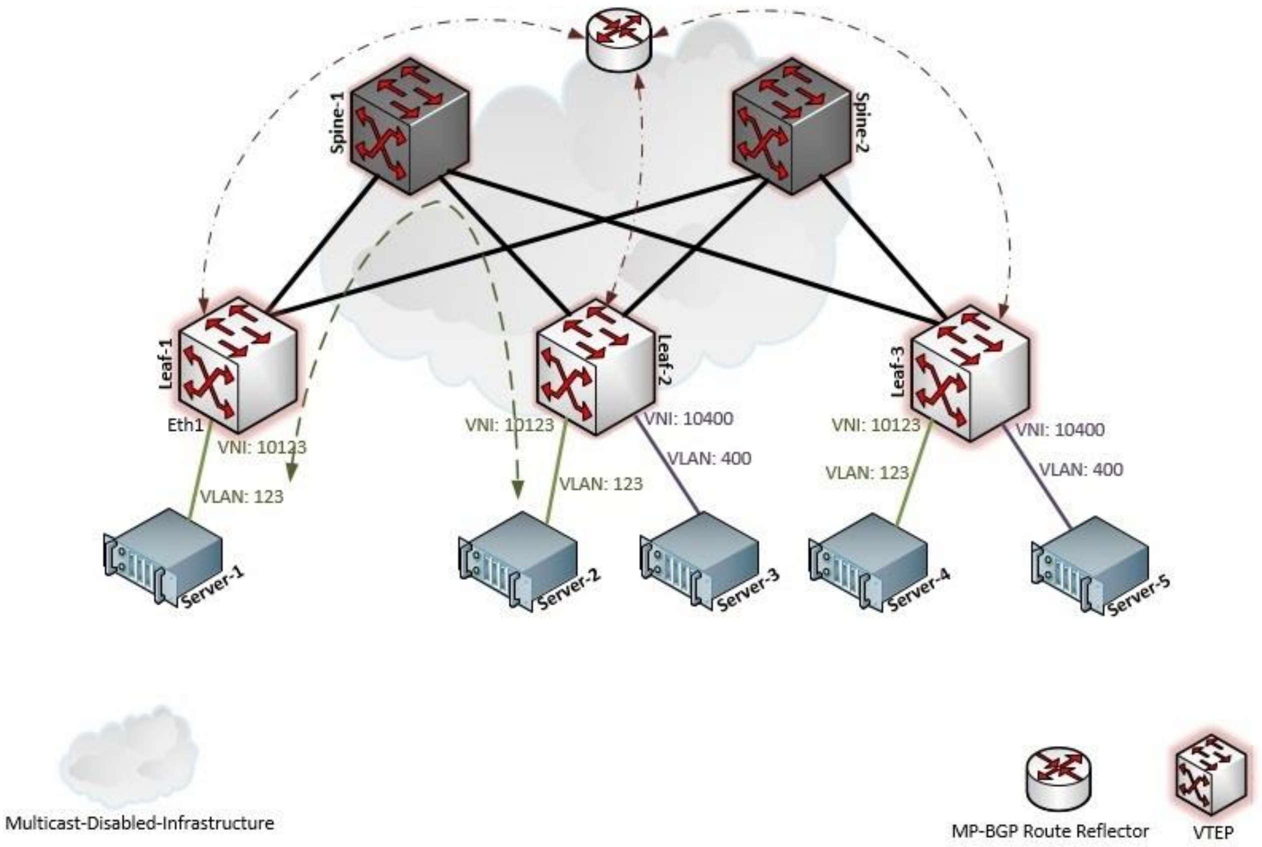

Figure 5: EVPN VXLAN

As in the controller-based VXLAN, there is no unknown unicast. The list of all communicating MACs is on each VTEP, and in case of an unknown MAC the local VTEP will direct the traffic via the default entry towards the VXLAN gateway.
Again for broadcast and multicast traffic, the head-end replication is always the solution.

BaaT VXLAN Control Plane Operation

BaaT achieves control plane operation via blockchain.

In this mode, the VXLAN VTEPs are also nodes of a public or private blockchain, as shown in Figure 6, that can span the public Internet. Note that there is always the option to replace blockchain with a DAPP (Decentralized Application) that runs over a blockchain-based platform, but this paper focuses on the blockchain option.

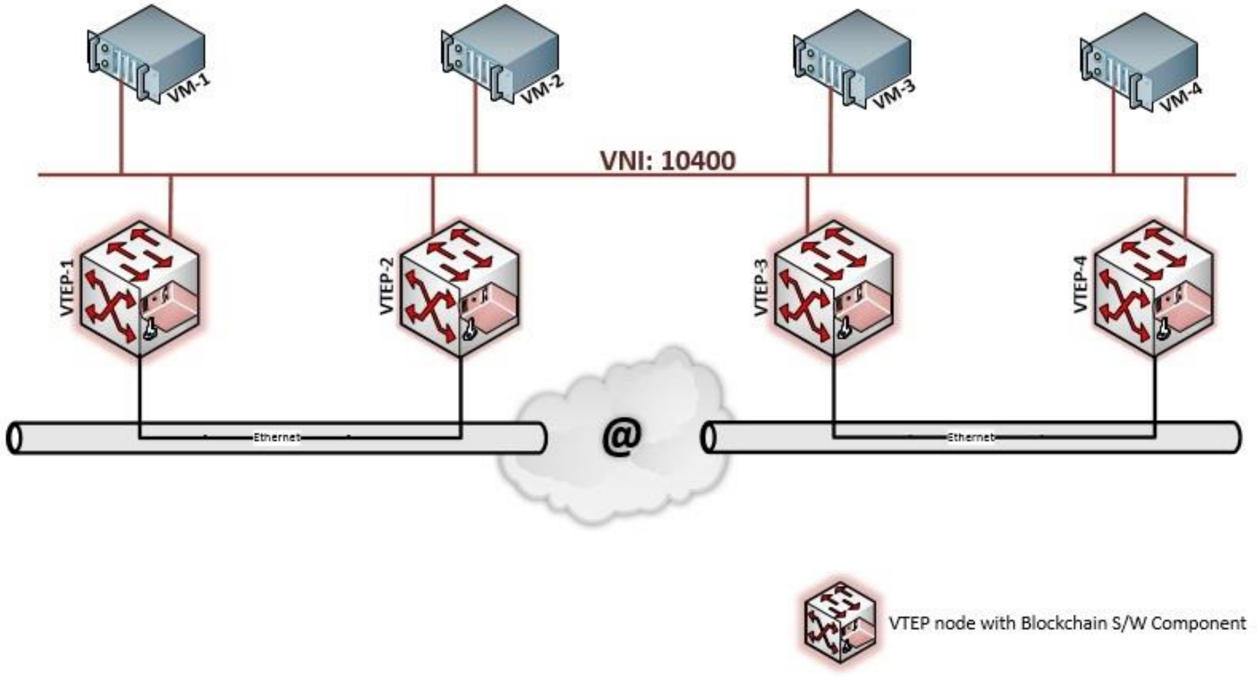

Figure 6: BaaT Global and Local

It's worth noting that a hybrid blockchain-- as opposed to a public blockchain. It is also less expensive than public blockchains, which usually require payment every time there is a change. And, the hybrid blockchain respects an organization's privacy by allowing the organization to freely set its rules and consensus in a way that services its policy.

The local MAC learning technique is the same as with any other VXLAN operation: The VTEPs learn the local MAC addresses via their local ports, as shown in Figure 7, and then the addresses are advertised/published as reachable through their VTEP IPs over the blockchain using transactions that are packed into proper blocks.

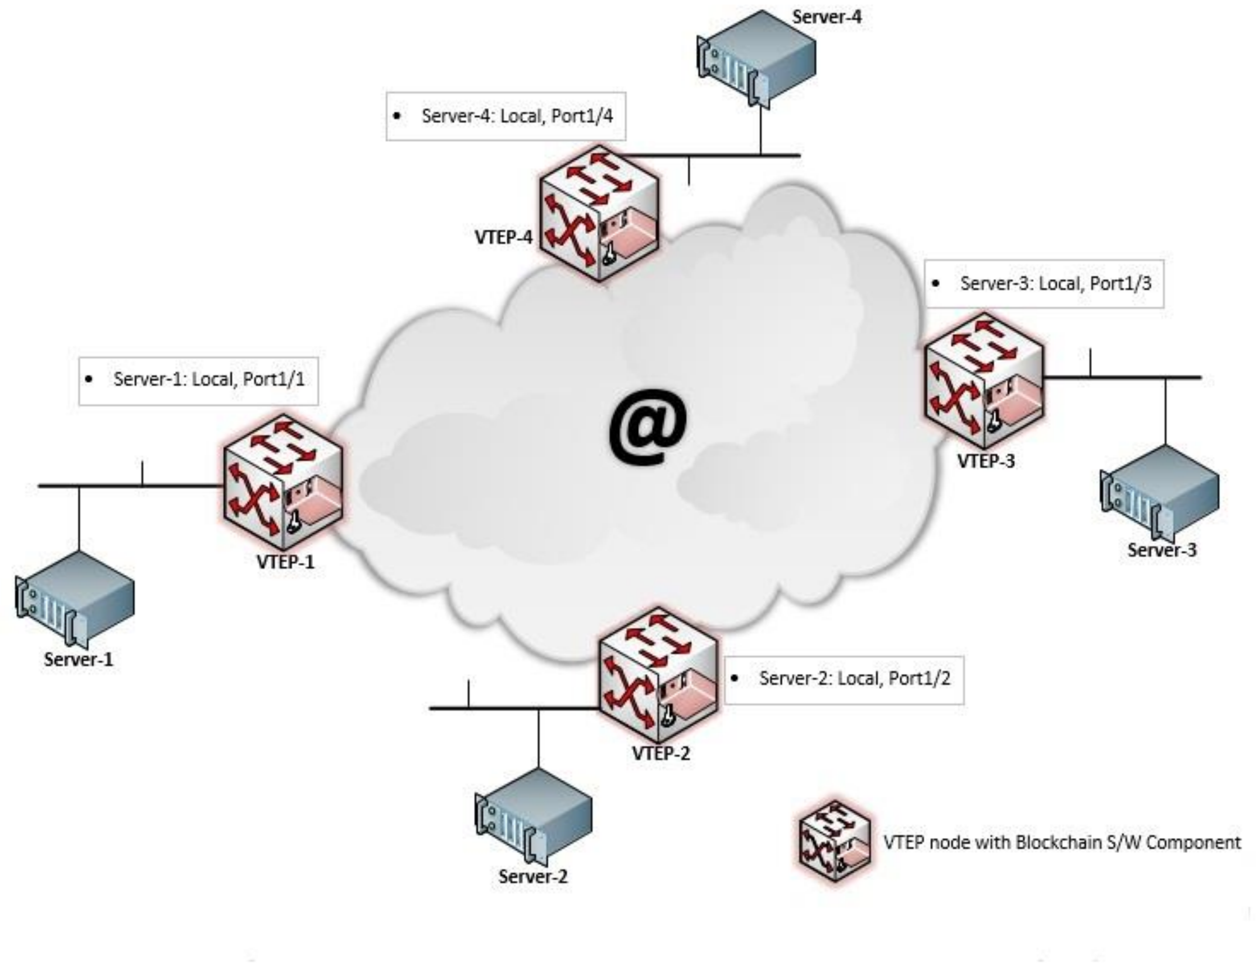

Figure 7: BaaT Initial State

Recall from Section 6.5,“Data Storage and Retrieval,”that an example was shown of successfully publishing a stream of hexadecimal data over the blockchain and we clarified that this stream is not random but meaningful.

Any online converter tool can be used to convert the hexadecimal information into text. In the example, we published:

‘637573746f6d65722031353520766e6964203130343030’

Converting hexadecimal to text, we find:

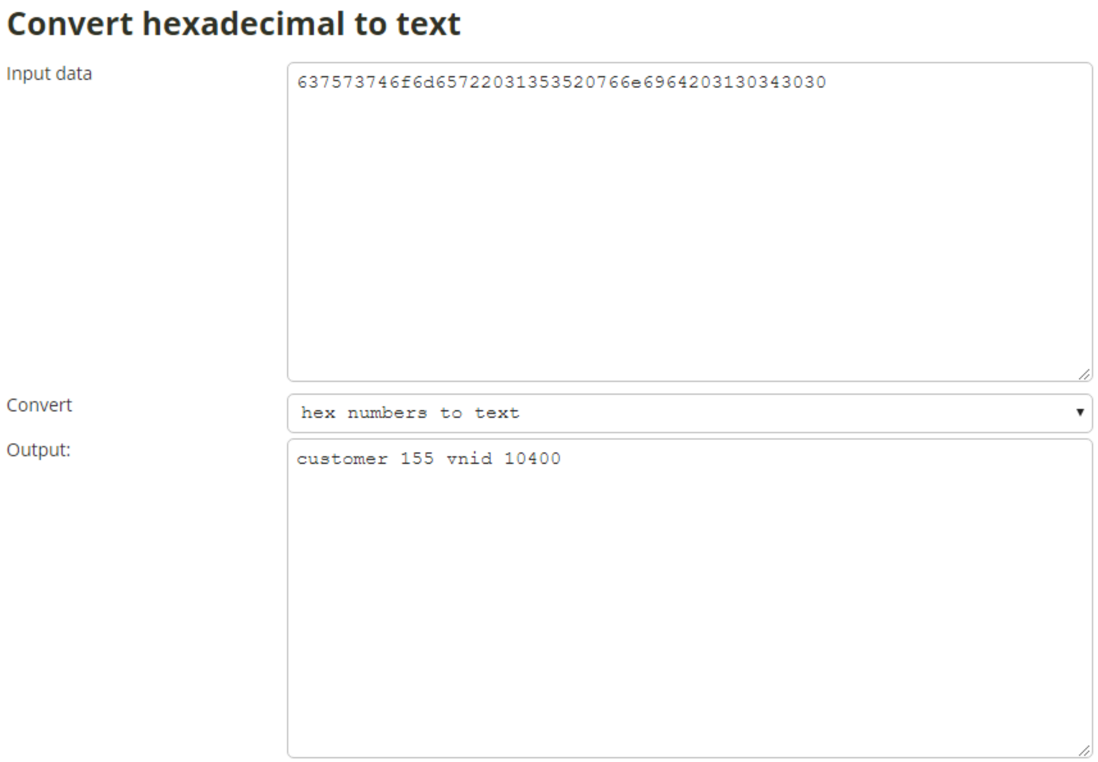

Convert hexadecimal to text

Input data

```
637573746f6d65722031353520766e6964203130343030
```

Convert: hex numbers to text

Output:

```
customer 155 vnid 10400
```

Figure 8: Converting Hexadecimal to Text

The output text is:

‘**customer 1555 vnid 10400**’.

This simple output message can be used to carry information related to the BaaT operation for a given customer.

From another node, we published:

‘**637573746f6d6572203135353520766e6964203132353030**’.

This means: ‘**customer 1555 vnid 12500**’ -- another VNID for that same customer.

The actual BaaT implementation will not be as straightforward as what has been illustrated. Instead, we use complex encryption techniques to properly encrypt the sensitive information sent over the blockchain. Also, the published messages won’t be as simple as the example here: The messages must include all information pertaining to a specific client for the proper BAAT operation. Consider this use case: Customer #1555 LAN segment is connected to VTEP-1 (in Figure 8), VTEP-1 needs to participate in VXLAN with VNID 10123. VTEP-1 just learned the MAC address (00-14-22-01-23-45**)** from one of the locally attached servers belonging to Customer#1555 (Server-1 attached to Port 1/1).

The VTEP-1 IP address is 10.1.1.178, and this is the IP address that other VTEPs need to use to reach VTEP-1.

So the message that will be published from that VTEP over the blockchain is typically a MAC-to-VTEP mapping message that also includes the Customer ID as well as the VNID.

In this use case, the message will be:

‘**customer 1555 vnid 10123 mac address 00-14-22-01-23-45 VTEP 10.1.1.178**’

Its hexadecimal format that will actually be published over the blockchain is:

‘**637573746f6d6572203135353520766e6964203130313233206d61632061646472657373203030302d31342d32322d30312d32332d343520565445502031302e312e312e313738**’

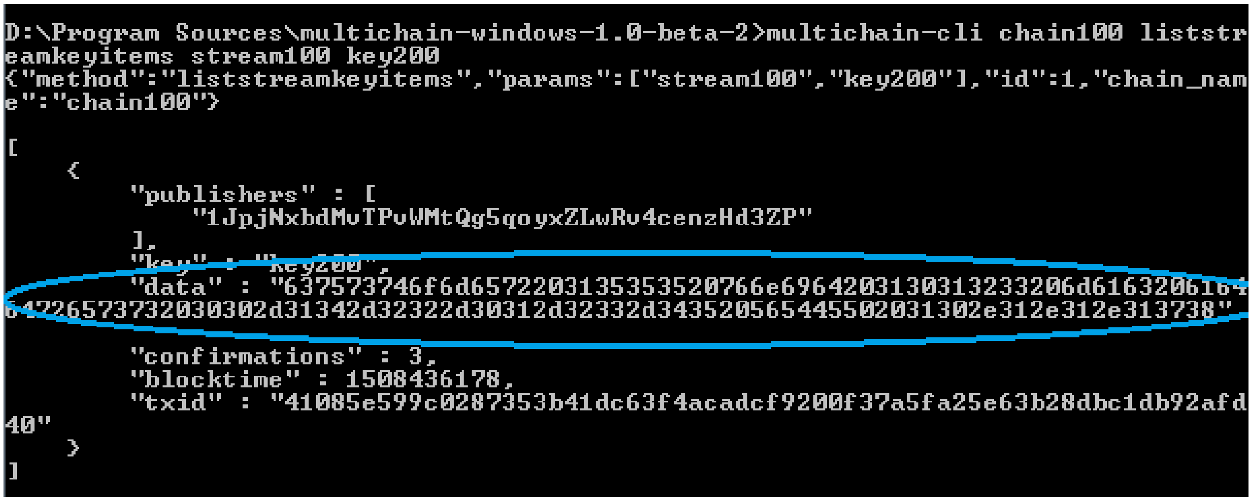

```
D:\Program Sources\multichain-windows-1.0-beta-2>multichain-cli chain100 liststr
eamkeyitems stream100 key200
{"method":"liststreamkeyitems","params":["stream100","key200"],"id":1,"chain_nam
e":"chain100"}

[
    {
        "publishers" : [
            "1JpjNxbdMvTPvWMtQg5qoyxZLwRv4cenzHd3ZP"
        ],
        "key" : "key200",
        "data" : "637573746f6d6572203135353520766e6964203130313233206d616320614
64726573732030302d31342d32322d30312d32332d343520565445502031302e312e312e313738",
        "confirmations" : 3,
        "blocktime" : 1508436178,
        "txid" : "41085e599c0287353b41dc63f4acadcf9200f37a5fa25e63b28dbc1db92afd
40"
    }
]
```

This message can be seen by all VTEPs participating in the blockchain but only those VTEPs that are interested in Customer ID 1555 and VXLAN VNID 10123 will use this message, translate it, and add its content to their local copy of the MAC-to-VTEP mappings.

Because the different MAC-to-VTEP mappings are distributed over the blockchain to all participating nodes/VTEPs, as the final state in Figure9:

- No data-plane learning is required for unknown unicast MAC addresses.
- No IP multicast underlay required. This is why BaaT can span beyond the boundary of a data center or cloud to the public Internet.
- Because of the distributed nature of blockchain, no significant delay is expected between the different nodes.
- For the broadcast and multicast traffic, the head-end replication is always the solution as in other control-plane-based VXLAN modes.

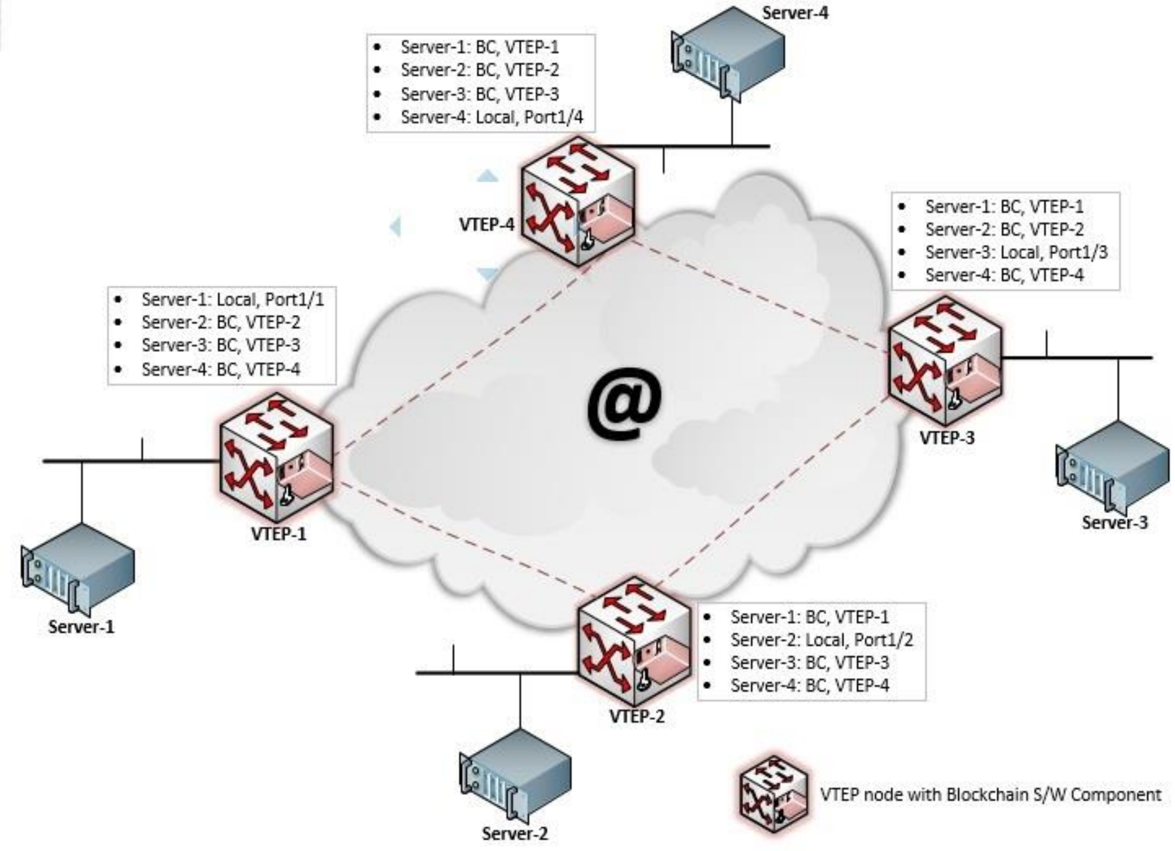

Figure 9: BaaT Final State

Hacking and MITM Concerns

In the case of Encrypted Network Overlay version of BaaT using the public Internet it is important to understand how security concerns such as hacking and MITM attacks are addressed.

<u>External Hacking Attacks are thwarted by:</u>

- **Endpoints of Tunnel are INVISIBLE** = (created via SHA-256 Cryptography and impenetrable). Private IP addresses and the associated VLANs are encrypted. The only way to "participate" in a conversation is to be an authenticated node on the blockchain – via verified MAC address database. All Public IP addresses are "white-listed" and circuits are hardened with all well-known and unverified TCP/UDP ports denied.
- **Full Encryption of Data in Motion** = via High Assurance AES-256 Encryption Key Management – and Quantum Resistant Algorithms based on Post-Quantum Cryptography Standardization.

<u>Internal Hacking Attacks</u> (including physical compromise)
Q. What Damage Could Be Done? A. ZERO Compromise.
BaaT circuit would "break" (keys would be flushed if box is tampered with - Physical device security for

secure storage of key materials with tamper detection and response circuitry.
Just as with External Hacking Attacks - it would take fifty supercomputers that could check a billion billion (1018) AES keys per second (if such a device could ever be made) and would, in theory, require about 3×10^{51} years to exhaust the 256-bit key space to break through the BaaT encryption method.

Man-in-the-Middle Attacks are thwarted with BaaT by:

- **Best practice Encryption Key Management** = random, unique cryptographic keys that are computed in the factory or when the system first boots up.
- **High Assurance (Robust) Encryption** = authenticated end-to-end encryption, where there is no point in the network, or link where un-encrypted data is accessible.

With High Assurance encryption solutions that provide true end-to-end encryption security, there is no weak point. Even if a router or switch (as is frequently the case) is found to have weak-points or vulnerabilities to attacks, High Assurance encryption ensures the data passing through those vulnerable devices is safe.