OpenCryptoTrust

Blockchain For Modern Telecommunications

www.OpenCT.io

2019 Technical Documentation for:

BaaT

Blockchain-as-a-Transport

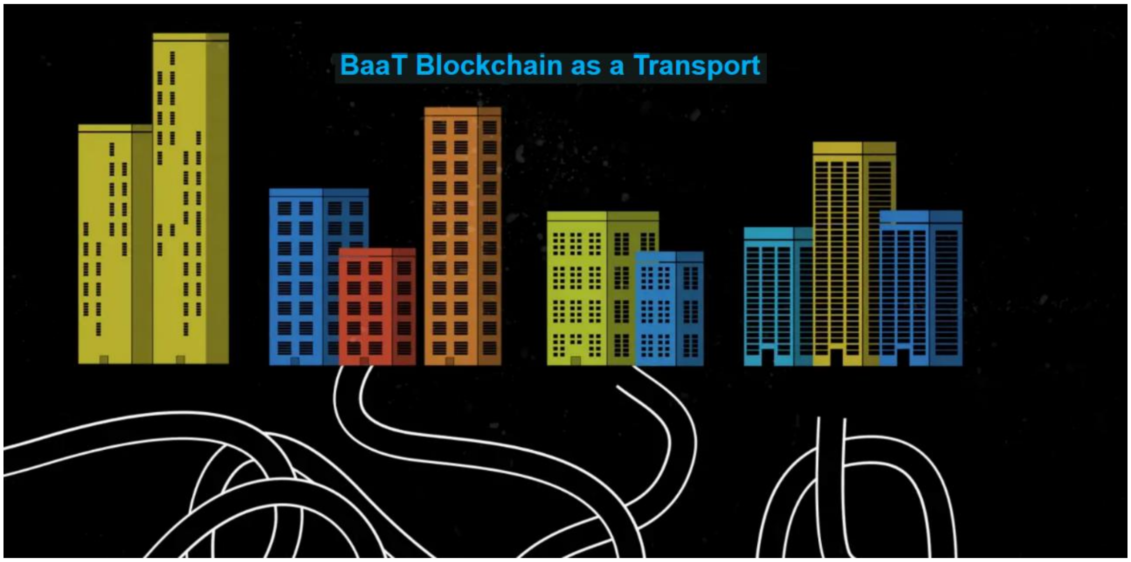

An Overview of BaaT – Blockchain-as-a-Transport

*Described is **Blockchain-as-a-Transport (BaaT)** networking architecture. BaaT connects geographically dispersed Layer 2 (L2) islands over any available infrastructure, including the public Internet. BaaT securely supports all kinds of network traffic including; Unicast, Multicast, and Broadcast.*

BaaT leverages blockchain to create an architecture that can connect geographically dispersed Layer 2 (L2) islands over any available infrastructure, including the public Internet.

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.

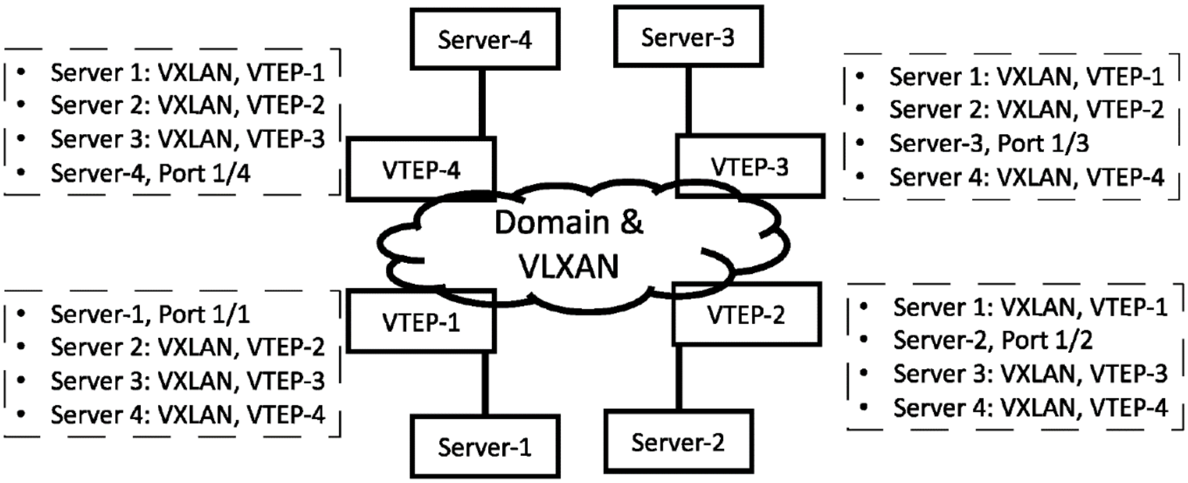

Figure 1.

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

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

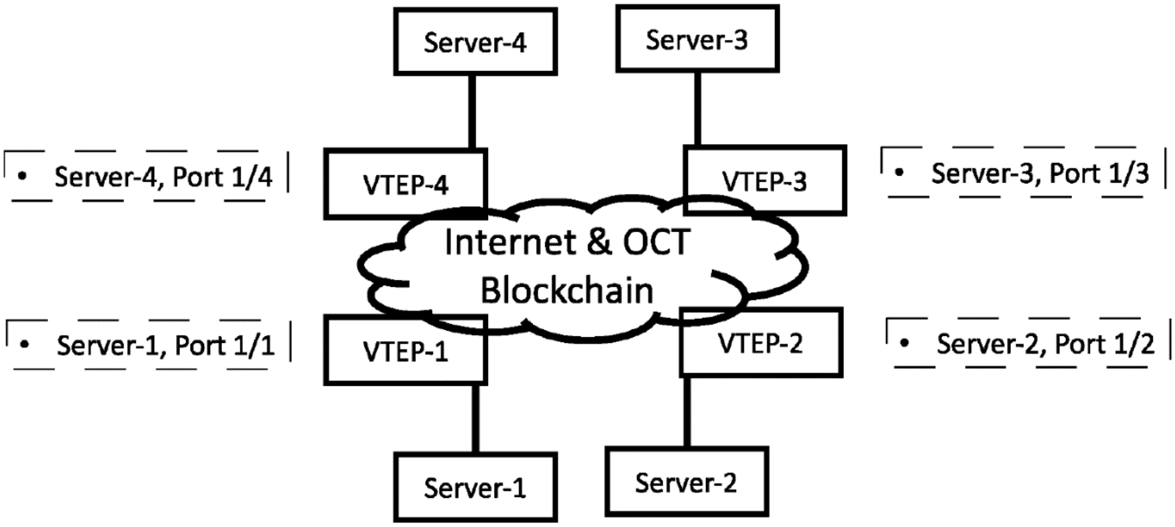

Figure 2.

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

See Figure 9. Hexadecimal Text Transmission

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.

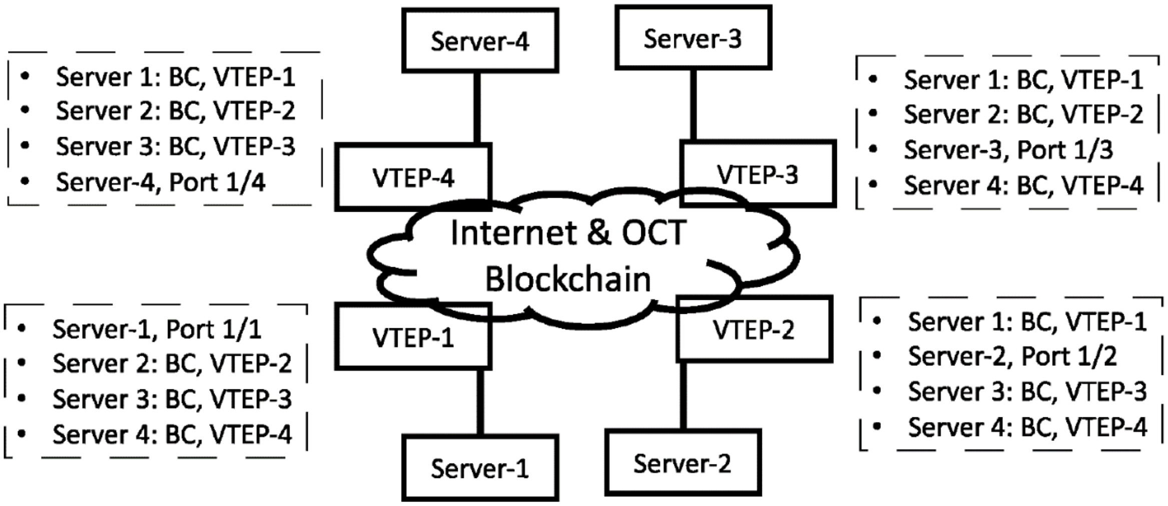

Figure 3.

1.1 BaaT Unique Security Strength

In our globally connected world, cyberattacks have become an ever-increasing risk for enterprises and governments. Data transported over public networks must be protected to meet privacy and integrity requirements and to ensure compliance with existing and emerging regulation.

BaaT has been developed in compliance with the strongest security standards such as the US Federal Information Processing Standard (FIPS) 140-2, issued by the National Institute of Standards and Technology (NIST).

Unlike traditional VPN based solutions BaaT leverages two (2) distinct security methodologies:

A) Transport Layer Security via Blockchain **SHA-256 – Secure Hash Algorithm**
B) Data Packet Encryption via **AES-256 – Advanced Encryption Standard**

The combination of these two methods ensures that neither the endpoints (source/destination) or the data payload can be penetrated or compromised.

To explain BaaT's security strength (uniquely using SHA-256 for TLS) it is helpful to explain the power of Hash Functions.

A Hash Function is a mathematical function which returns a very random digest or a number. Hash function applied on the same message will generate the same output. There is no way to get the original message by somehow decrypting the digest.
Even a small change in the message changes the digest drastically.
One such example of a hash function used widely is the SHA256 function which returns a 256 bit number.

SHA256("Hey")

10101110111011010110111100010101010010001011111101101011010111

SHA256("HEy")

000000101010101011111000100010100001101010101011011100110

A 256 bit pattern can represent 2^{256} different possible messages. To give you an idea of how difficult it is to extract the original message : Breaking a symmetric 256-bit key by brute force requires 2^{128} times more computational power than a 128-bit key. Fifty supercomputers that could check a billion billion (1018) AES keys per second (if such a device could ever be made) would, in theory, require about 3×10^{51} years to exhaust the 256-bit key space.

1.2 BaaT Data Encryption

BaaT data encryption functions with line-rate performance - 0.36 us (microsecond) for lowest latency and highest throughput. It's built on an enhanced version of MACsec with additional tunneling overlay so that no infrastructure needs to be replaced. BaaT provides corrective controls as part of a complete set of Carrier Ethernet CFM and OAM functions for service monitoring and testing.

Specified by IEE, 802.1AE MACsec is a method for protecting the confidentiality and integrity of Ethernet connections. This standard specifies a technique for securing links between nodes, decrypting incoming and encrypting outgoing traffic.

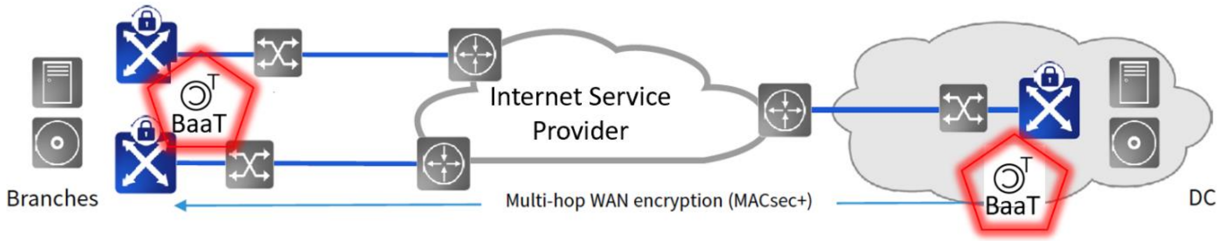

Highlights of BaaT's advanced data encryption features include:

- Hardware-based MACsec encryption with IEEE 802.1AE standard-based frame format.
- Advanced encryption standard IEEE 802.1AEbn-2011-complaint GCM-AES-256 NIST-approved cyphers.
- Physical device security for secure storage of key materials with tamper detection and response circuitry.
- Minimum overhead (24 bytes) to maximize throughput and minimize delay.
- Secure 1GE and 10Gbit interfaces with demarcation devices and edge compute nodes.
- Password-authenticated Diffie-Hellman key exchange protocol securing against man-in-the-middle attacks.

Understanding the advantages of BaaT requires first understanding the operation of VXLAN. In the following subsections, an operational overview of traditional VXLAN leads to a technical description of BaaT operation.

1.3 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.*VXLANextends 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.

1.4 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

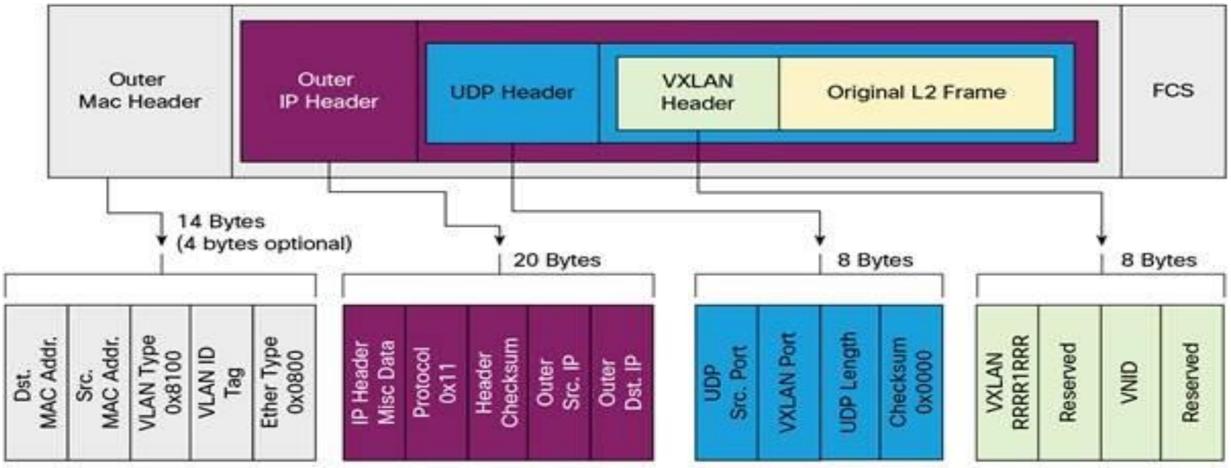

Figure 1.1: VXLAN Encapsulation

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.

1.5 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.

1.5.1 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.

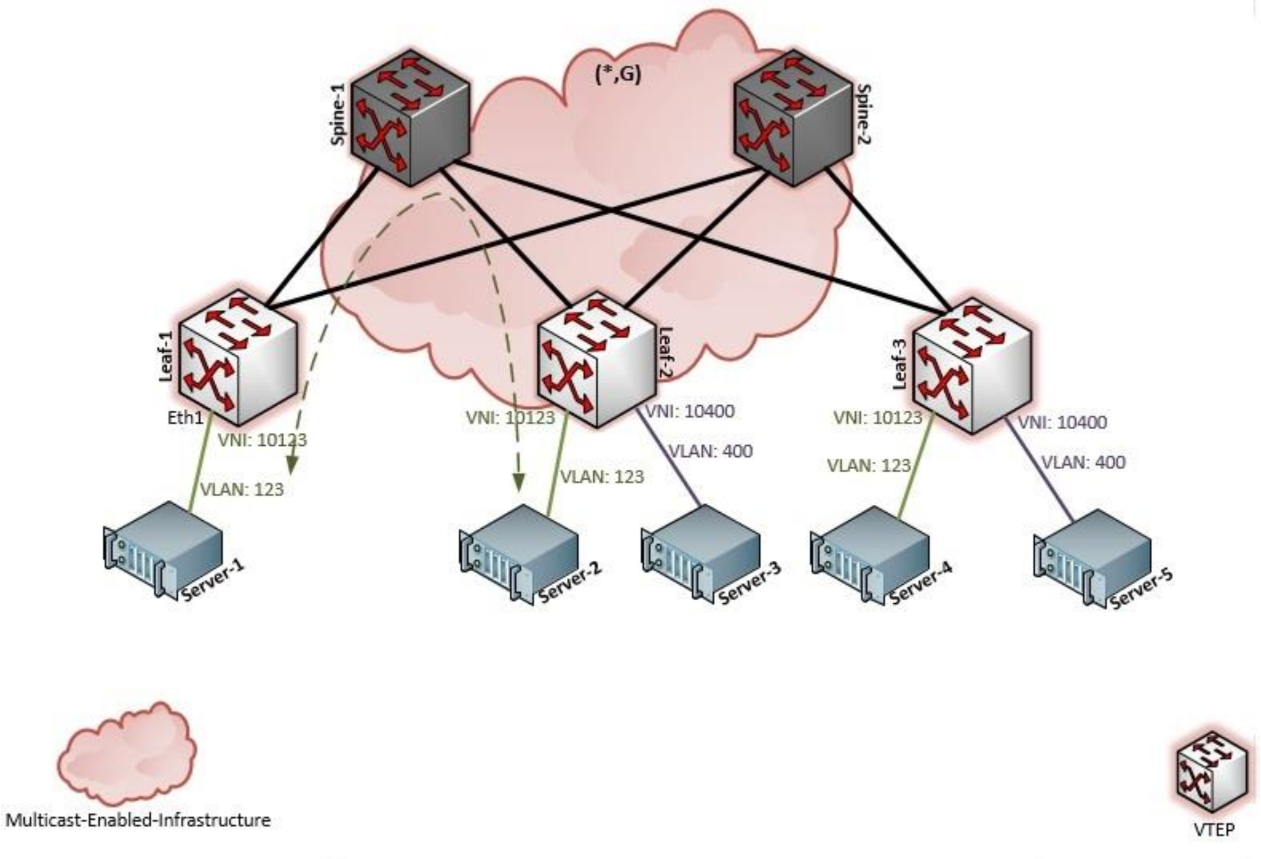

Figure 2.1: Control-Plane-Less Multicast VXLAN

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.

1.5.2 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.

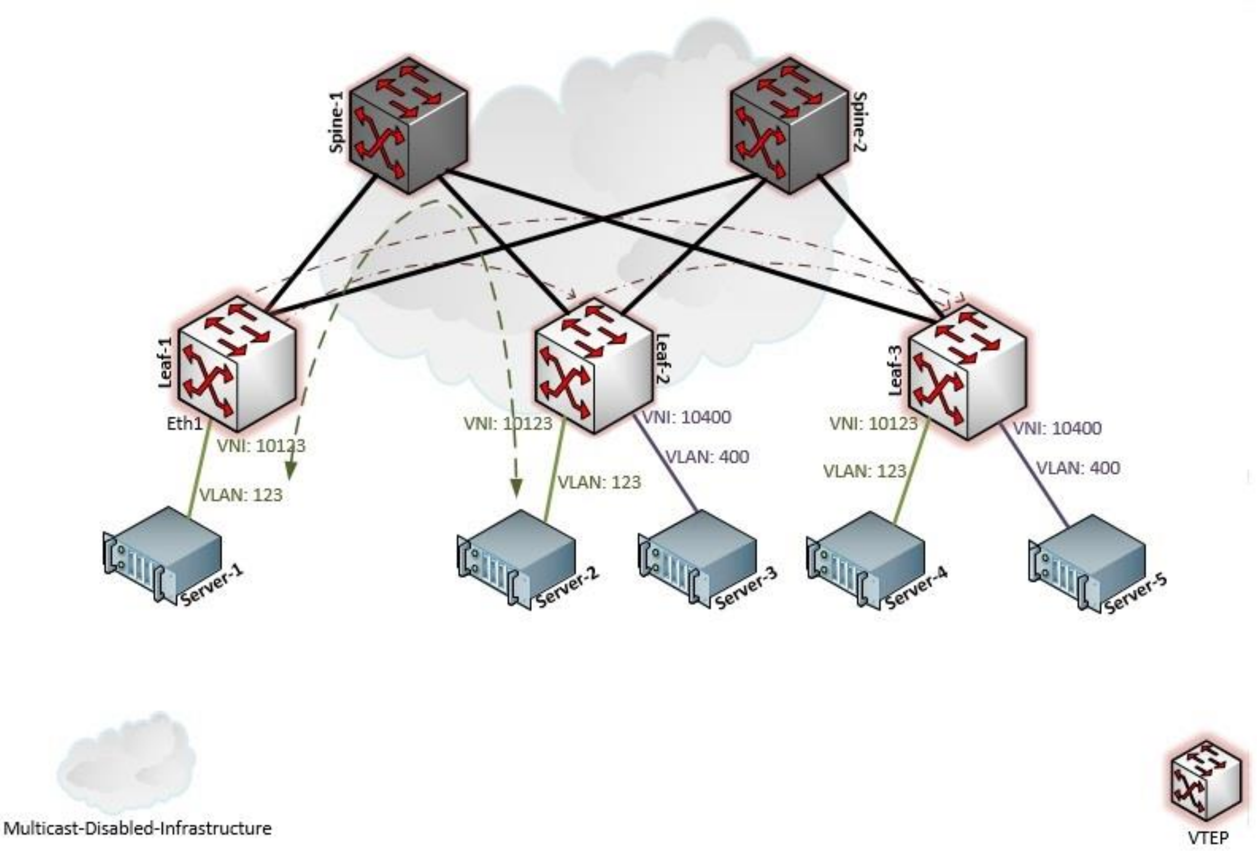

Figure 3: Control-Plane-Less Unicast VXLAN

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.

1.5.3 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.

1.5.4 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.

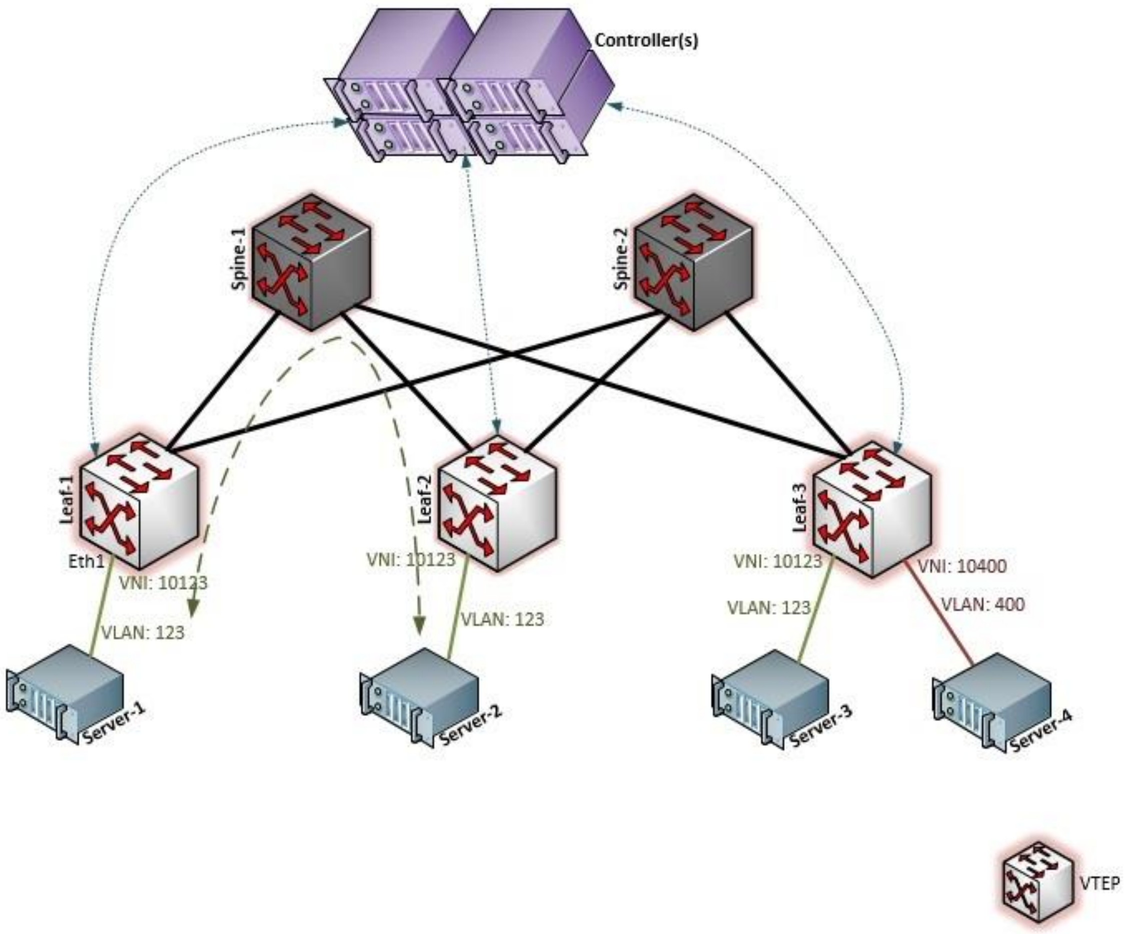

Figure 4: Controller-Based VXLAN

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.

1.5.5 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.

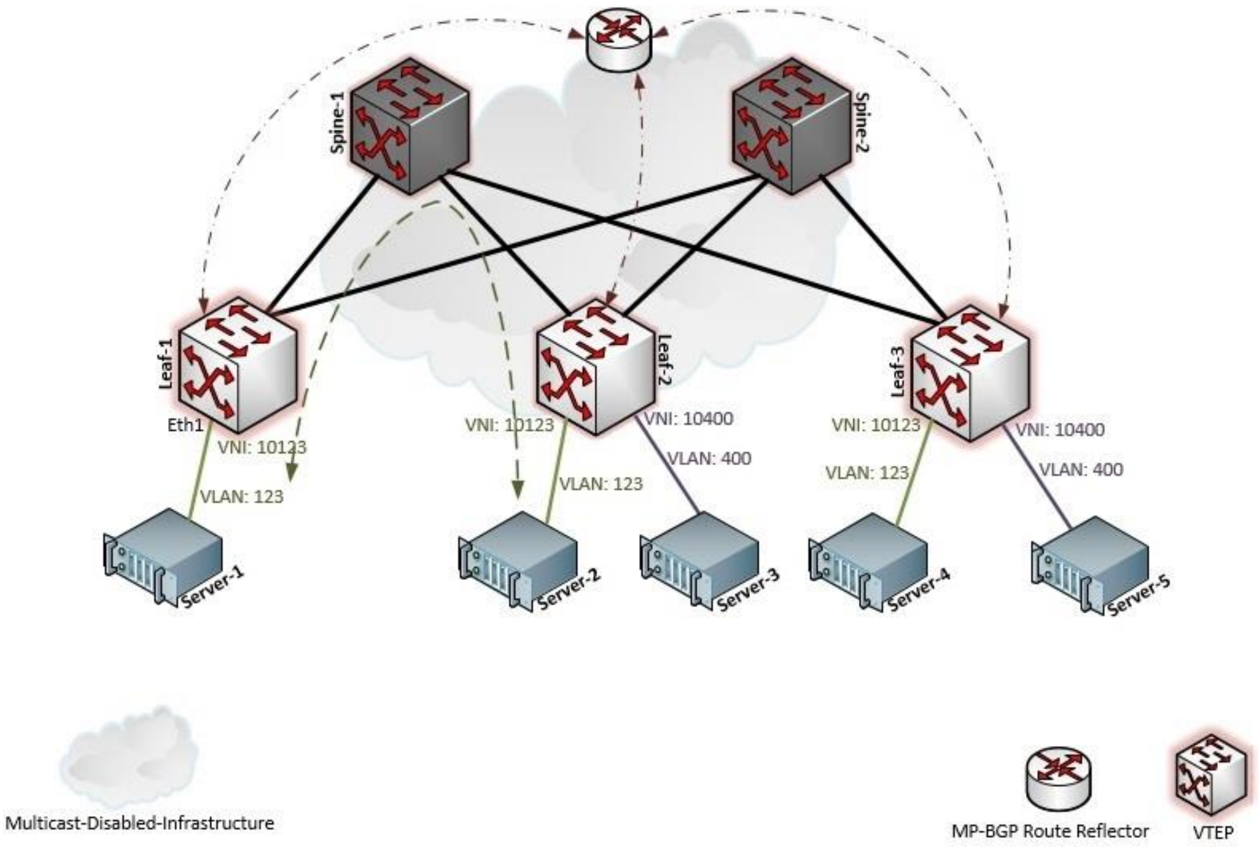

Figure 5: EVPN VXLAN

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.

1.6 BaaT or Blockchain VXLAN

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.

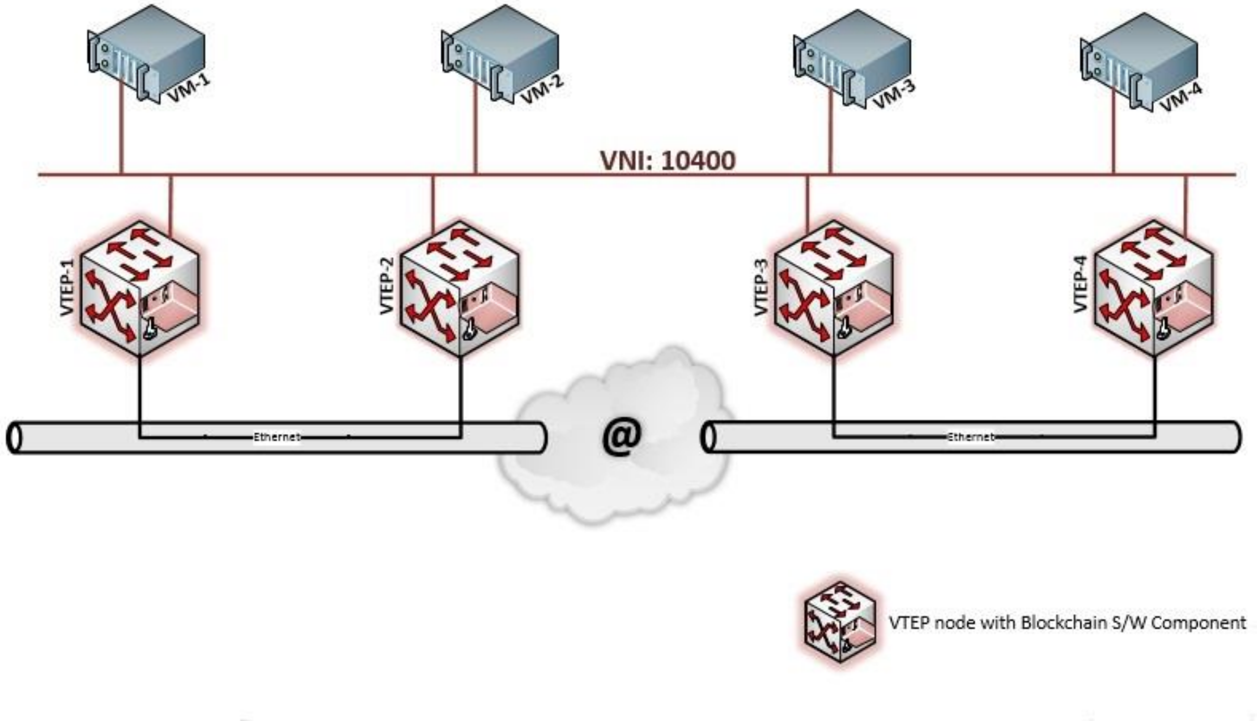

Figure 6: BaaT Global and Local

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.

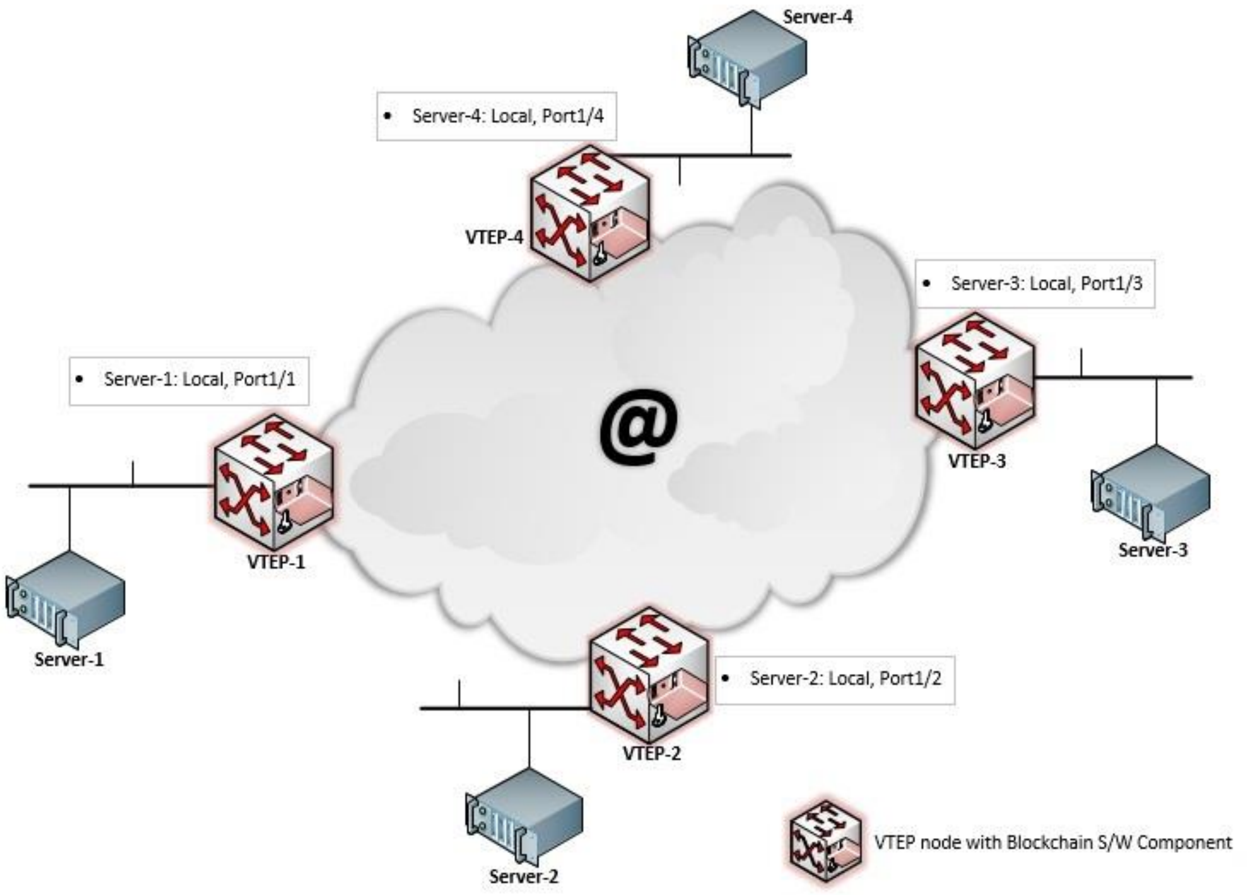

Figure 7: BaaT Initial State

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:

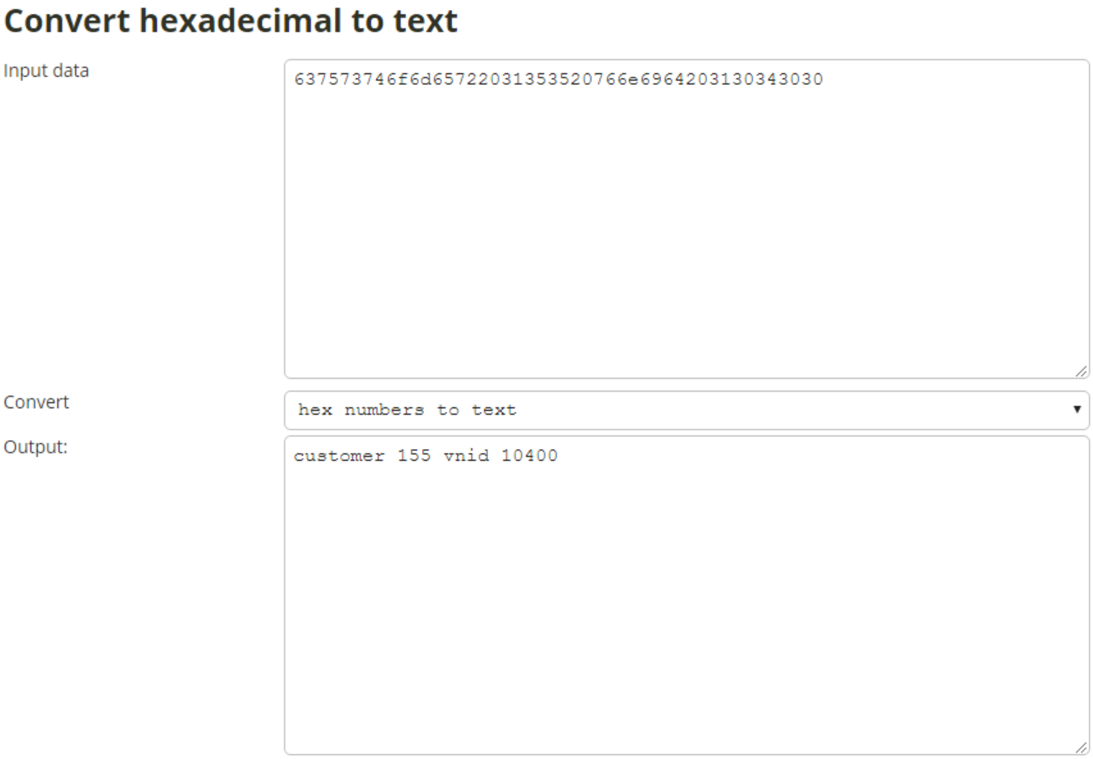

Convert hexadecimal to text

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:

‘**637573746f6d6572203135353520766e6964203130313233206d616320616464726573732030302d31342d32322d30312d32332d343520565445502031302e312e312e313738**’

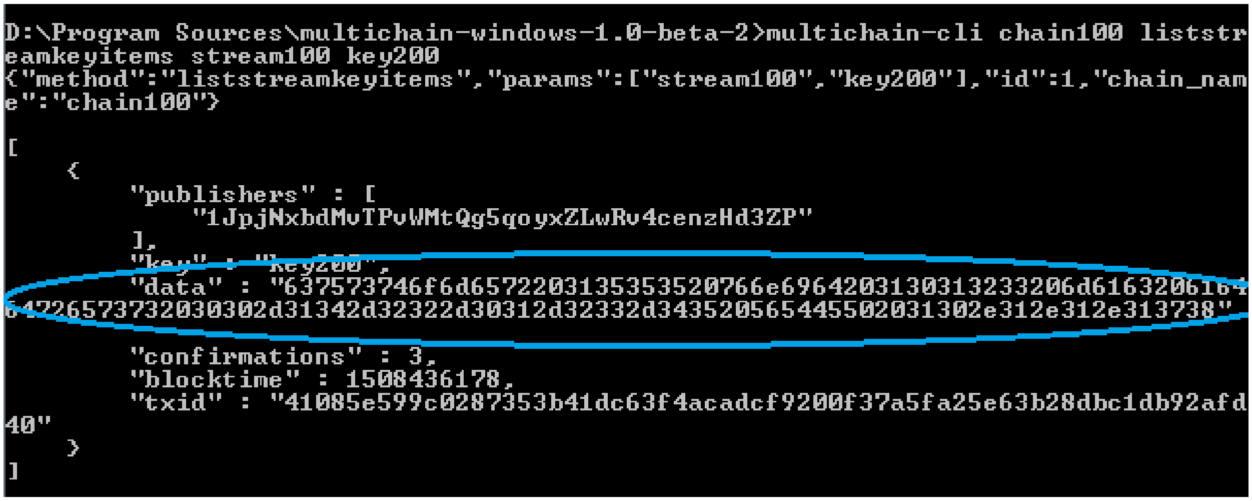

```
D:\Program Sources\multichain-windows-1.0-beta-2>multichain-cli chain100 liststr
eamkeyitems stream100 key200
{"method":"liststreamkeyitems","params":["stream100","key200"],"id":1,"chain_nam
e":"chain100"}

[
    {
        "publishers" : [
            "1JpjNxbdMvTPvWMtQg5qoyxZLwRv4cenzHd3ZP"
        ],
        "key" : "key200",
        "data" : "637573746f6d6572203135353520766e6964203130313233206d6163206164
6472657373203030302d31342d32322d30312d32332d343520565445502031302e312e312e313738",
        "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.

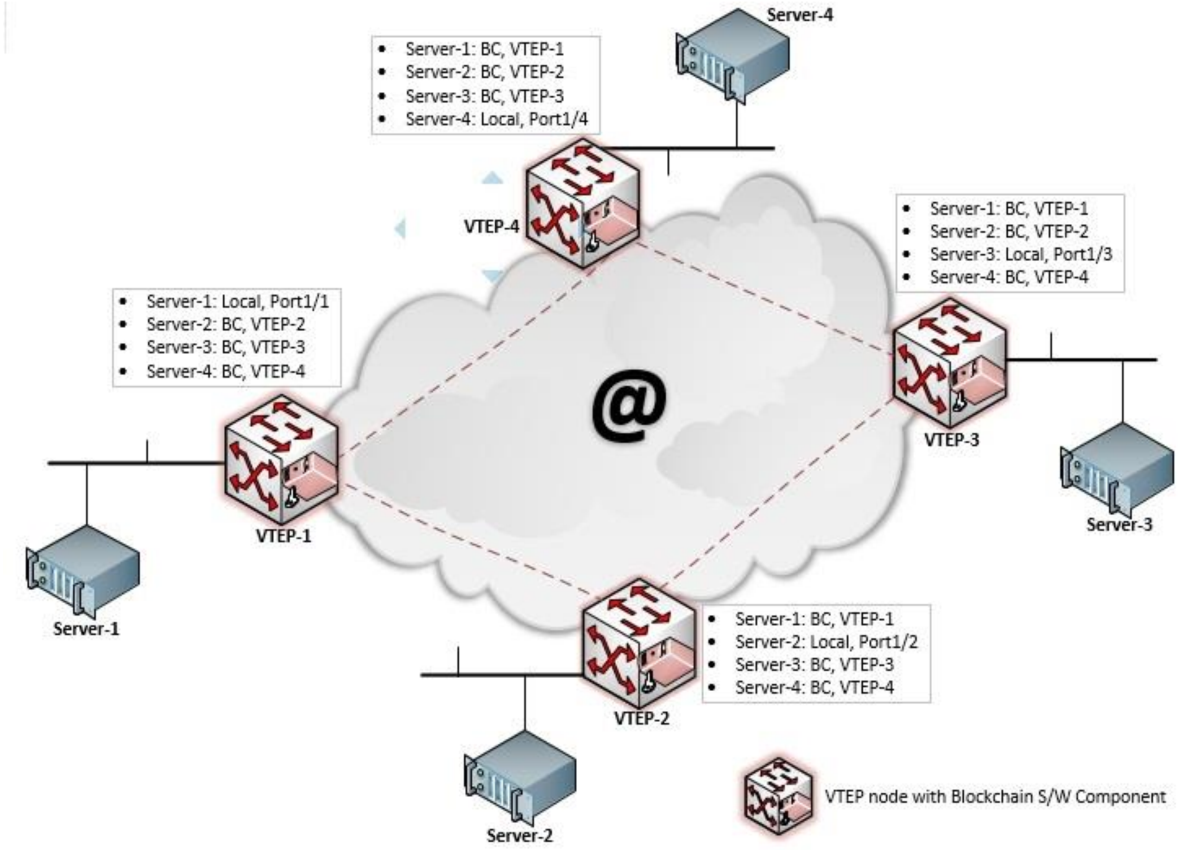

Figure 9: BaaT Final State

1.7 BaaT Value Proposition

- BaaT is an appealing WAN transport option compared to traditional WAN technologies like VPN, MPLS, or expensive dedicated links because of its seamless operation across the public Internet.
- BaaT is an advanced L2VPN solution for enterprises, governments, and telco/service providers, with which the organization can leverage the public Internet for their WAN traffic so that they don't need to share traffic with their upstream providers as in the case of MPLS VPN service or even modern SD-WAN.
- Unlike other tunneling techniques, BaaT is built to operate in a multipoint fashion. Its signaling is done separately via the blockchain and it has no scalability issues.
- BaaT can operate over any IP transport network including the global public Internet because no multicast underlay network is required.
- No unknown unicast entry is to be found on any VTEPs; it's either a unicast MAC address match, advertised over the blockchain, or a default entry toward the VXLAN gateway.
- Multicast and broadcast traffic is handled via the head-end replication on the source VTEP to all other known VTEPs in the same VXLAN.The list of VTEPs is known - and always updated - over the blockchain.

To conclude, BaaT is the Layer 2 transport service of choice for all businesses requiring absolute security and efficiency.