



UNIVERSITY OF OREGON

Let SDN Be Your Eyes: Secure Forensics in Data Center Networks

Adam Bates

University of Oregon

Kevin Butler

University of Oregon

Andreas Haeberlen

University of Pennsylvania

Micah Sherr

Georgetown University

Wenchao Zhou

Georgetown University

SENT'14, San Diego, CA, USA


23 February, 2014

Secure Forensics



UNIVERSITY
OF OREGON

- Investigating the possibility of a security breach is extremely difficult.
- When suspicious events may be malicious or benign, finding an explanation for the event can be a tedious, manual task.
- Due to the possibility of advanced persistent threats, fast detection and investigation of anomalies is essential.



When parts of your network have been compromised, who can you trust to provide answers about the attack??

Let SDN Be Your Eyes



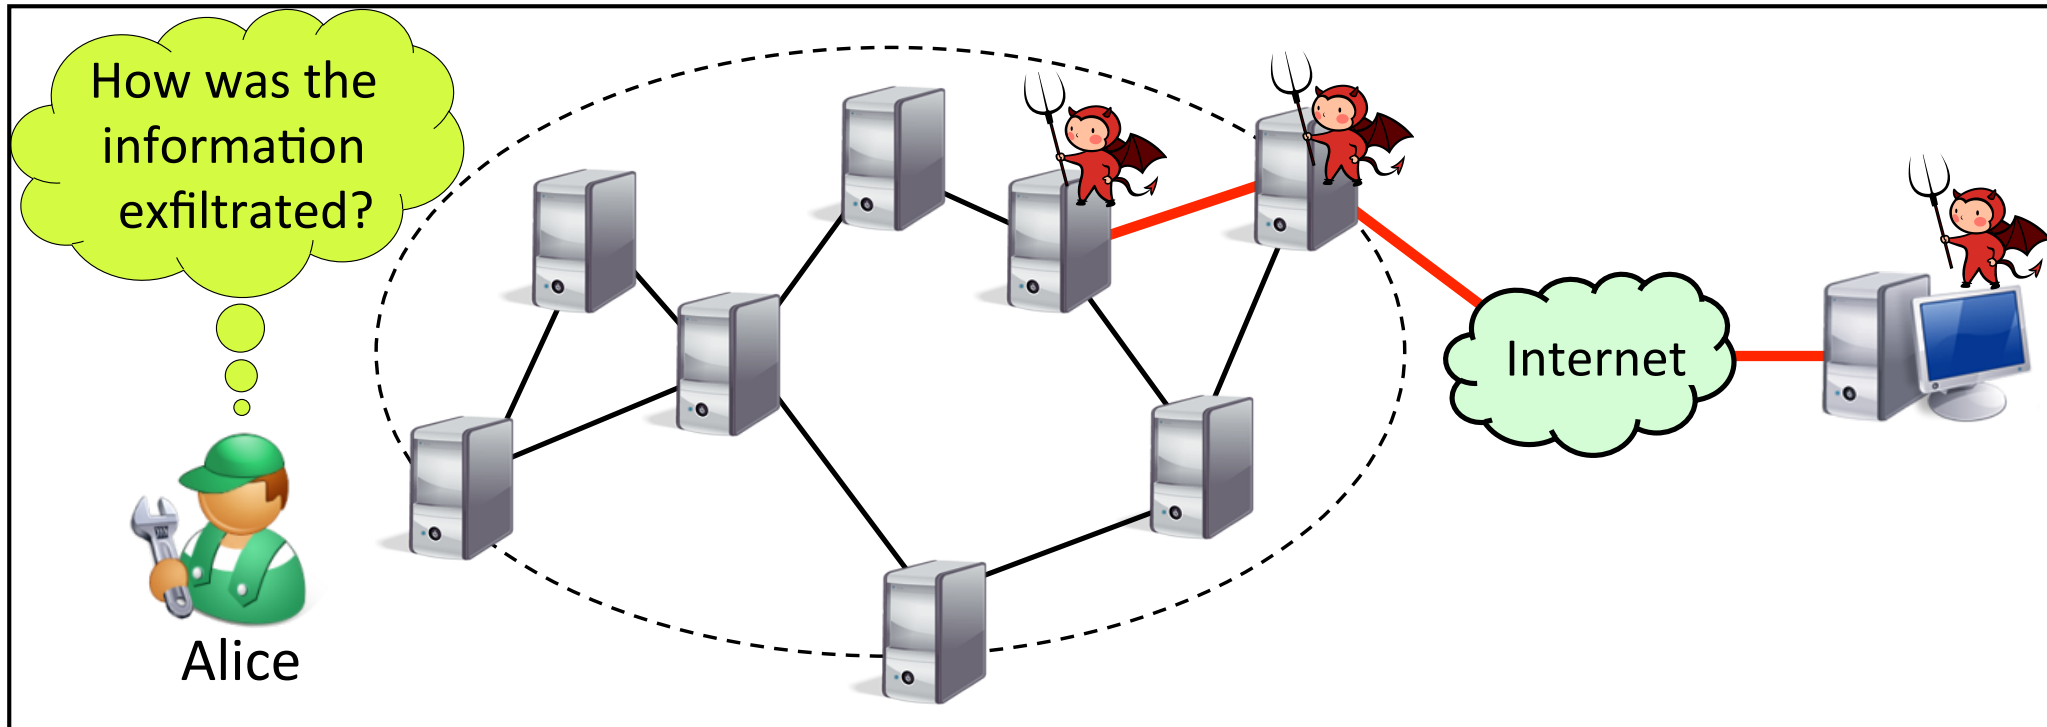
- We propose that Software Defined Networking (SDN) can be used to bootstrap trust in network forensics for data centers.
- We present an SDN-based *network provenance* system that extends into the network itself, creating a secure monitoring layer for all network activity.
- Our system possesses the ability to:
 - **Detect Covert Communication**
 - **Detect Equivocation**
 - **Detect Missing Forensic Records**

Unexpected Behavior?



UNIVERSITY
OF OREGON


Question: *Who can Alice trust for answers about the network?*



Unexpected Behavior?




Question: *Who can Alice trust for answers about the network?*



Answer: *Trust the available correct nodes in the network.*
(Secure Network Provenance [ZFN+2011])

- Secure Network Provenance (SNP) constructs a *provenance graph* based on system execution.
- Each node manages its own tamper evident log and follows an inter-node message commitment protocol.
- An administrator queries nodes' logs to reconstruct a provenance graph, detecting faulty nodes through finding inconsistencies and omissions.
- **Limitation**: SNP anchors trust in a critical mass of *correct* nodes; its view of the network shrinks as more nodes fail or are compromised.



- **Limitation:** SNP anchors trust in a critical mass of *correct* nodes; its view of the network shrinks as more nodes fail or are compromised.

Our Key Insight: Instead of anchoring trust at the host, we can use SDN to bootstrap trust in provenance-based network forensics.

- **Limitation:** SNP anchors trust in a critical mass of *correct* nodes; its view of the network shrinks as more nodes fail or are compromised.

Software-Defined Networking



- Programmable switches that facilitate decoupling of network's *data plane* from an abstract *control plane*.
- A Network Controller can instruct switches to handle flows in different ways. This is accomplished by pattern matching on network packet headers.
- SDN switches offer modest on-board functionality for packet processing, such as forwarding, dropping, flooding, header modification, etc.
- Higher-level network functionality is achieved by forwarding packets to “Middleboxes”, which can actually now be placed anywhere in the network.

Our system sets the following security goals:

- **Prevent Covert Communication:** eliminate all unmonitored paths for explicit communication between nodes within the network.
- **Detect Equivocation:** catch nodes that attempt to make inconsistent claims about their activities.
- **Response Availability:** In the event that message transcripts are missing, this should always be detectable by the system and the administrator.


Threat Model & Assumptions

- Nodes are subject to Byzantine Faults due to compromise or system failure.
- Faulty nodes may take actions collectively or individually to hide their presence from administrators.
- ALL communication in the network takes place over a network of SDN switches.
- SDN switch security, while important, is also not considered here. Switches are trusted.

What do we need from SDN?

- We are looking for a trustworthy *global observer* of network events.
- Since SDN switches offer constrained functionality, we introduce **Provenance Verification Points (PVP)**, a middlebox for forensic analysis of network packets.
- Switches force all flows through a PVP.
- Switches perform access control of network resources through communication with the PVP/Controller.

- Provide a verification layer to a host-level message commitment protocol [HKD07].
- All network traffic is duplicated at the switch and forwarded to both the recipient and a PVP.
- During forensic investigation, the administrator queries both the nodes and the PVP in order to get an accurate explanation for network activity.



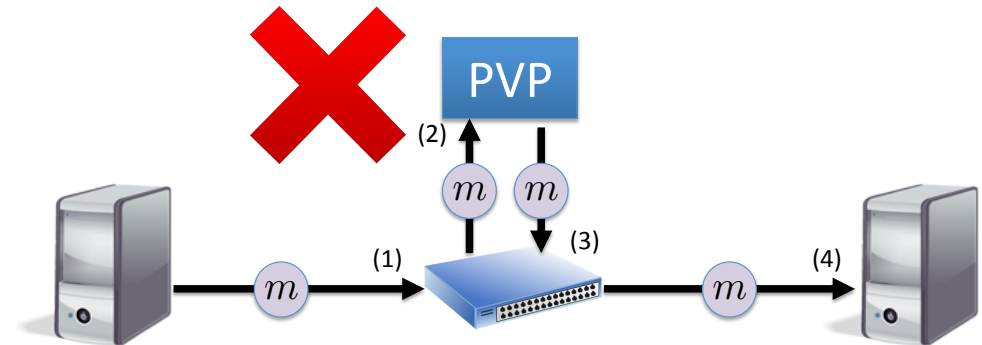
PVPs are a distributed set of middleboxes that observe and record node-to-node communications.

Design Consideration #1




Should the PVPs interpose on traffic, or receive mirrored traffic?

- Interposition imposes additional latency.
- Mirroring means that PVPs cannot actively detect or prevent exfiltration, but...
- If a protocol violation is detected, the PVP notifies the Network Controller to isolate the culprit.



Option 1: PVPs intercept all messages.




Option 2: PVPs receive copies of all messages.

Design Consideration #2

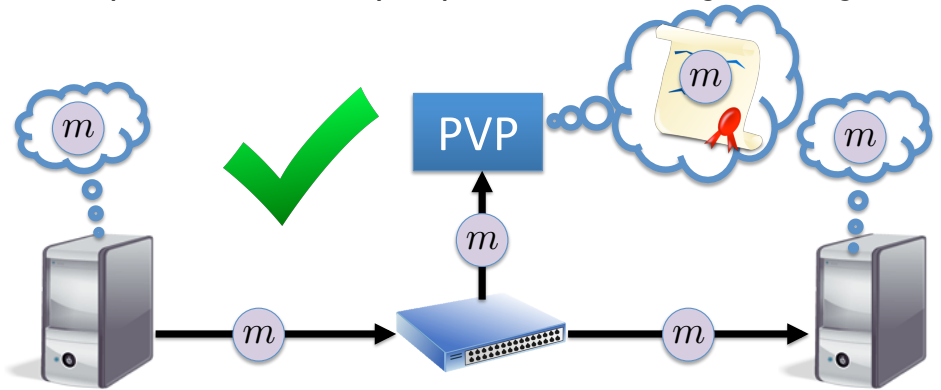


Should nodes participate in the network provenance protocol?

- *IF PVPs are more trustworthy, cutting out nodes would be more secure...*
- **But if PVPs do ALL of the work, it will be difficult to keep up with the network in real time...**
- **Instead, the PVP performs the minimum work needed to assure security goals.**




Option 1: PVP is solely responsible for message tracking.



Option 2: Nodes track messages sent and received, PVP maintains minimal proof of transmissions.

Message Commitment Protocol


Node A wishes to send a message m to Node B...



Node A wishes to send a message m to Node B...

Node A records m , a signature of his log at the current time, and a signature of his log at the immediate previous time.

PVP's Append-Only Log



Node A's Append-Only Log

$m, \sigma_A(\text{LOG}_{A,t}), \sigma_A(\text{LOG}_{A,t-1})$



Node A

Node B's Append-Only Log



Node B




Message Commitment Protocol


Node A wishes to send a message m to Node B...

Node A then sends the message and log commitments to Node B.


PVP's Append-Only Log




Node A's Append-Only Log



Node A



Node B's Append-Only Log

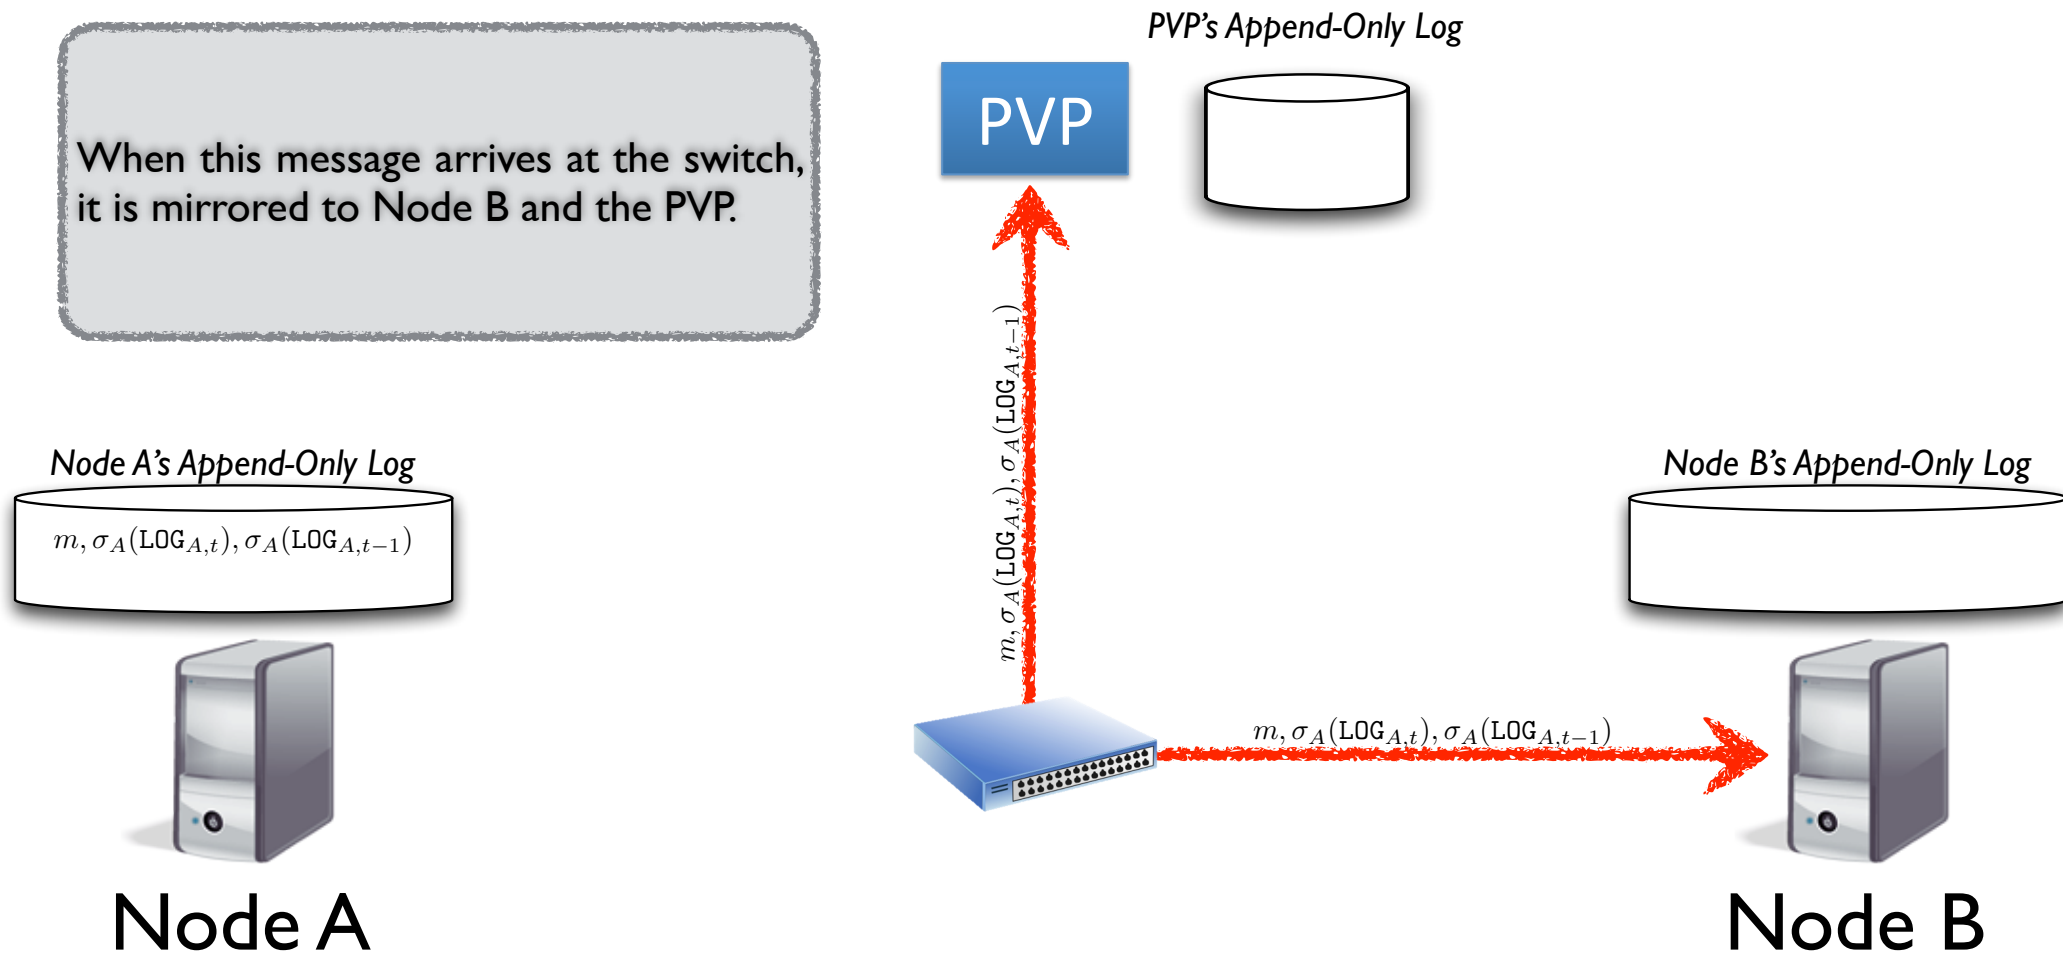


Node B

Message Commitment Protocol

Node A wishes to send a message m to Node B...


When this message arrives at the switch, it is mirrored to Node B and the PVP.




Node A wishes to send a message m to Node B...

After verifying the signatures, Node B logs the entire message from Node A. However, the PVP logs just the current signature (or “authenticator”).

PVP's Append-Only Log




Node A's Append-Only Log



Node A



Node B's Append-Only Log




Node B

Message Commitment Protocol


Node A wishes to send a message m to Node B...

Node B ACK's message m by following the same procedure, logging the message and then sending an ACK of m and log signatures back to Node A.

PVP's Append-Only Log




Node A's Append-Only Log



Node A


Node B's Append-Only Log



Node B



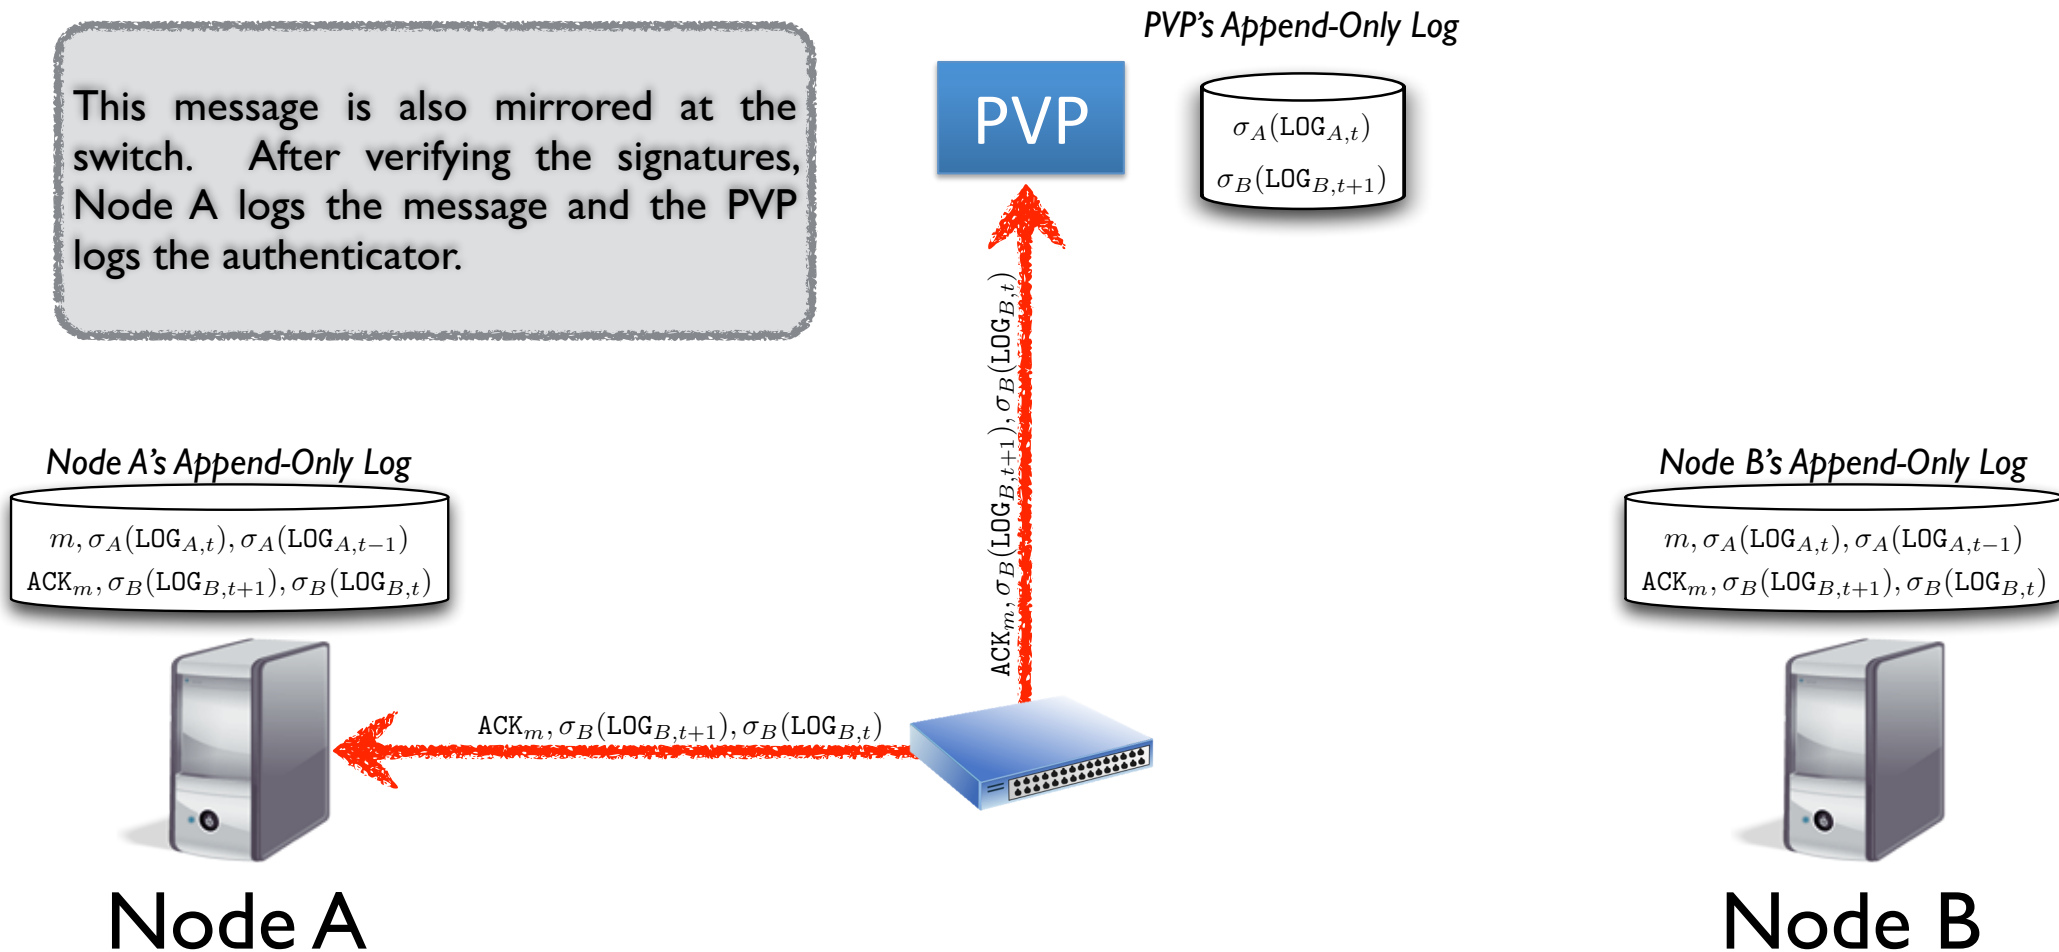
$\text{ACK}_m, \sigma_B(\text{LOG}_{B,t+1}), \sigma_B(\text{LOG}_{B,t})$






Message Commitment Protocol

Node A wishes to send a message m to Node B...


This message is also mirrored at the switch. After verifying the signatures, Node A logs the message and the PVP logs the authenticator.



Administrator asks “Why did m exist at time t ?”




Administrator asks “Why did m exist at time t ?”



Querying



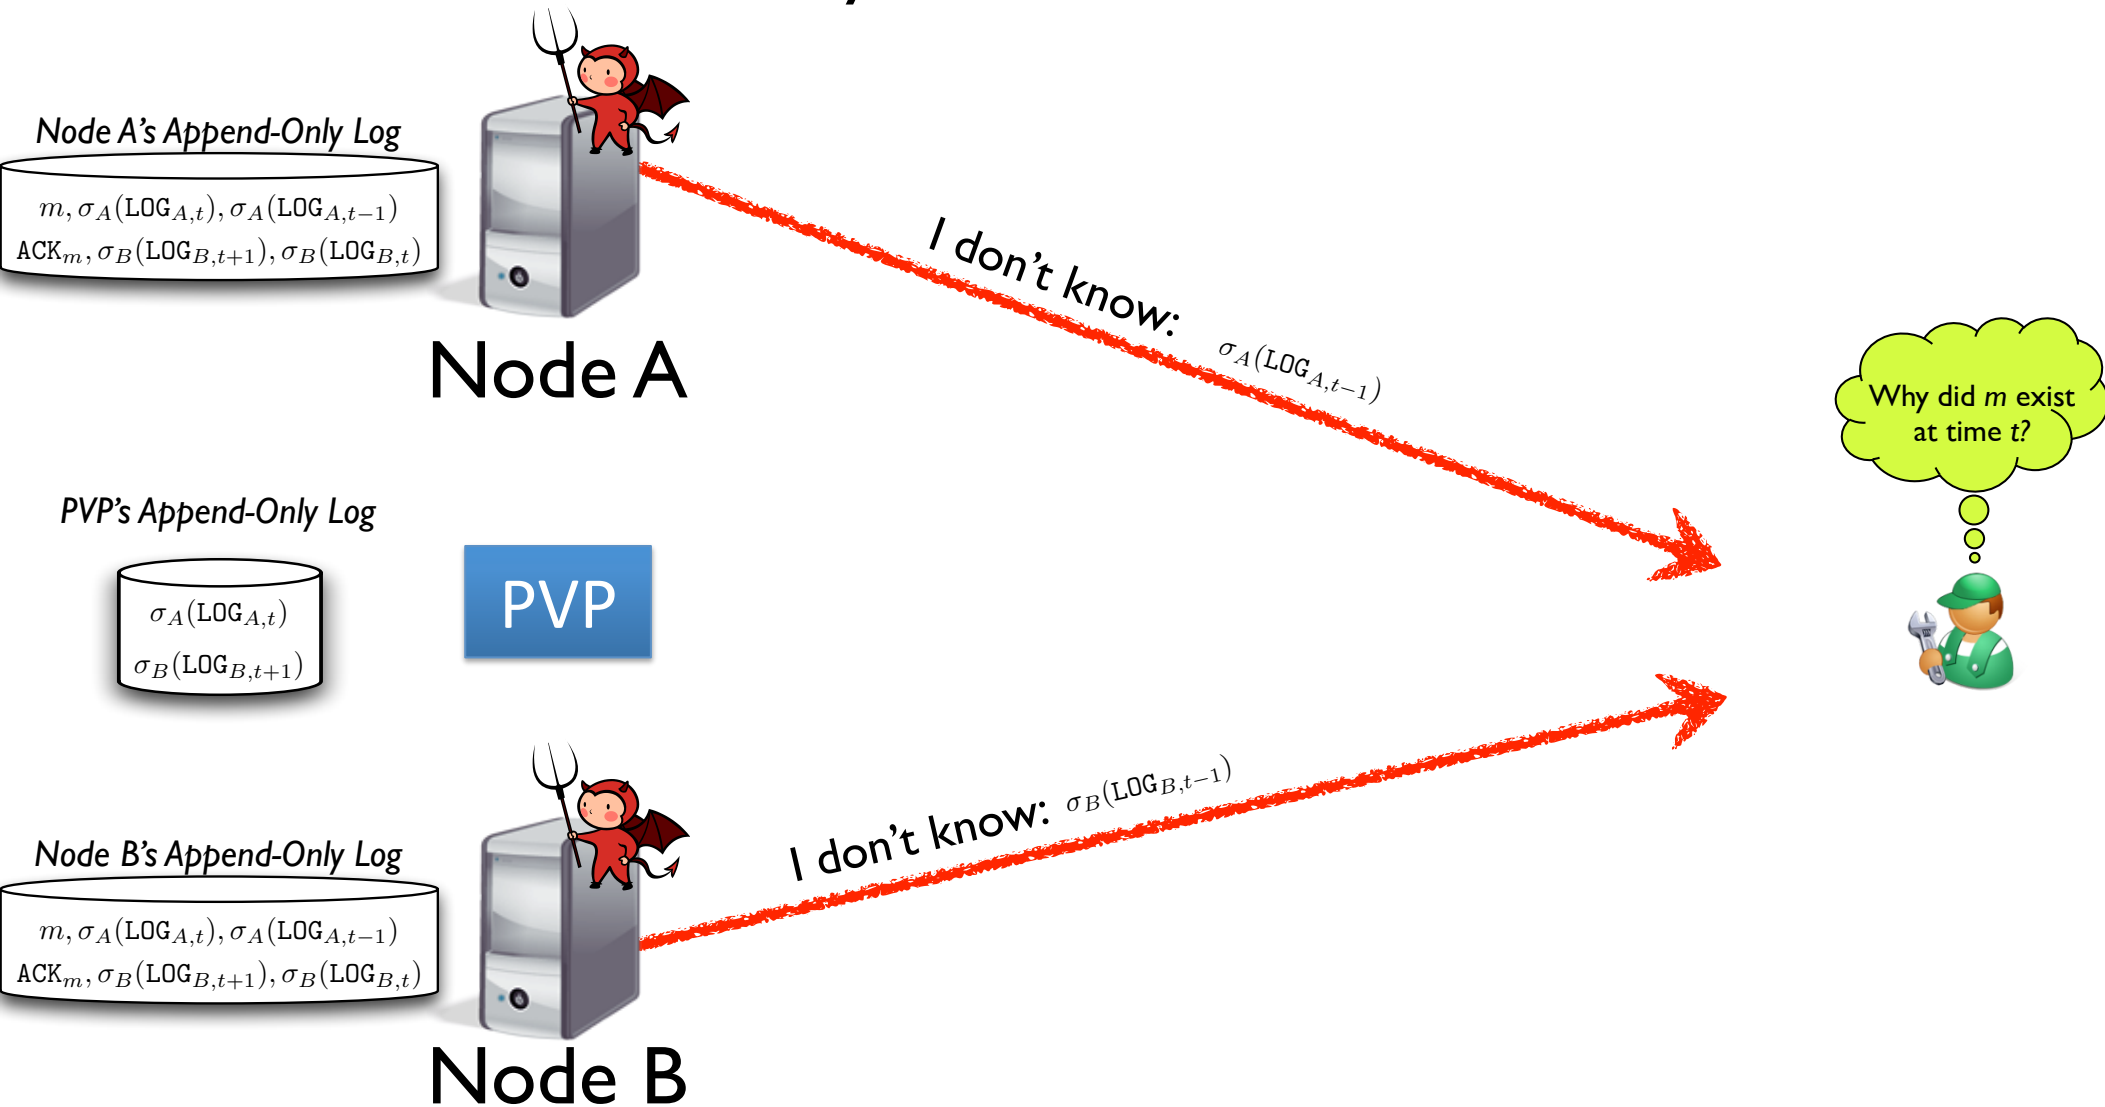
Administrator asks “Why did m exist at time t ?”




Querying




Administrator asks “Why did m exist at time t ?”



Administrator asks “Why did m exist at time t ?”



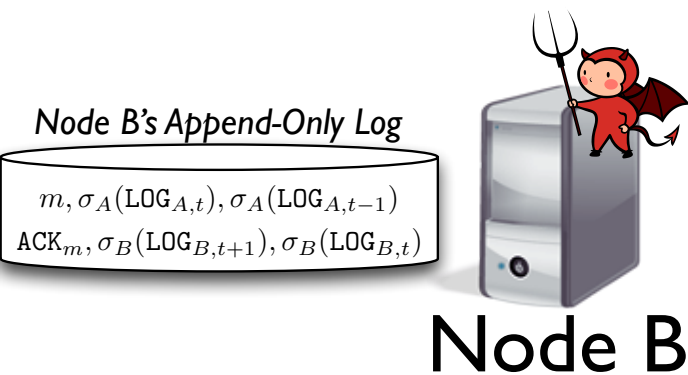

PVP's Append-Only Log



Here is proof of A and B's activity at time t :

$\sigma_A(\text{LOG}_{A,t})$

$\sigma_B(\text{LOG}_{B,t+1})$



Are PVPs Trustworthy?



- We do not have to explicitly trust the PVP. Because nodes keep a local log, they have proof to defend themselves with against a faulty PVP.
- If the PVP claims that a node sent an unauthenticated message, the node cannot defend itself. At this point, an administrator will need to resolve the conflict.
 - *We now know that either the PVP or the node is faulty!*
- In the presence of faulty PVPs, our security guarantees gracefully degrade to those of [ZFN+11].

- **Message Loss.** PVPs can recover by identifying retransmissions and polling the switch for flow statistics (This is an OpenFlow feature).
- **Timing Side Channels.** Incorporating timing information into our commitment protocol *may* permit Alice to later test for side channels.
- **Automated Forensics.** We believe that a policy-driven approach to network provenance would obviate the need for instrumenting applications to follow the message commitment protocol.

- We have shown that SDN can be used as a trust anchor to overcome limitations on previous network forensic systems.
- Using SDN, we have shown for the first time that reliable detection of covert communication between compromised hosts is possible.
- There are a variety of exciting opportunities and challenges in this area. We look forward to exploring them in future work.

Thank you



UNIVERSITY
OF OREGON

Any Questions?

Adam Bates

amb@cs.uoregon.edu

- Cost is the same as [ZFN+2011]
 - Network Overhead: Low as 0.2% (Hadoop MapReduce), high as 16.1% (Quagga BGP).
 - Computation: Quagga trials used 0.9% of one core without SNP, and just 5.4% of one core with SNP.
 - 365 Bytes Per Message.
- PVP Provisioning (Back-of-Napkin Estimates):
 - Can maintain all state within memory by hash chaining individual message authenticators.

Many parts of our system could be implemented through alternate means, so why is SDN necessary?

- Quickly isolate faulty nodes and redirect them to honeypots, ensuring that the protocol is followed.
- In Future Work, recovering from PVP failure will introduce dynamic routing challenges.
- In Future Work, recovering from message loss will require accurate and fine-grained flow-level statistics.