Trusted Disk Loading in the Emulab Network Testbed

Cody Cutler, Eric Eide, Mike Hibler, Rob Ricci

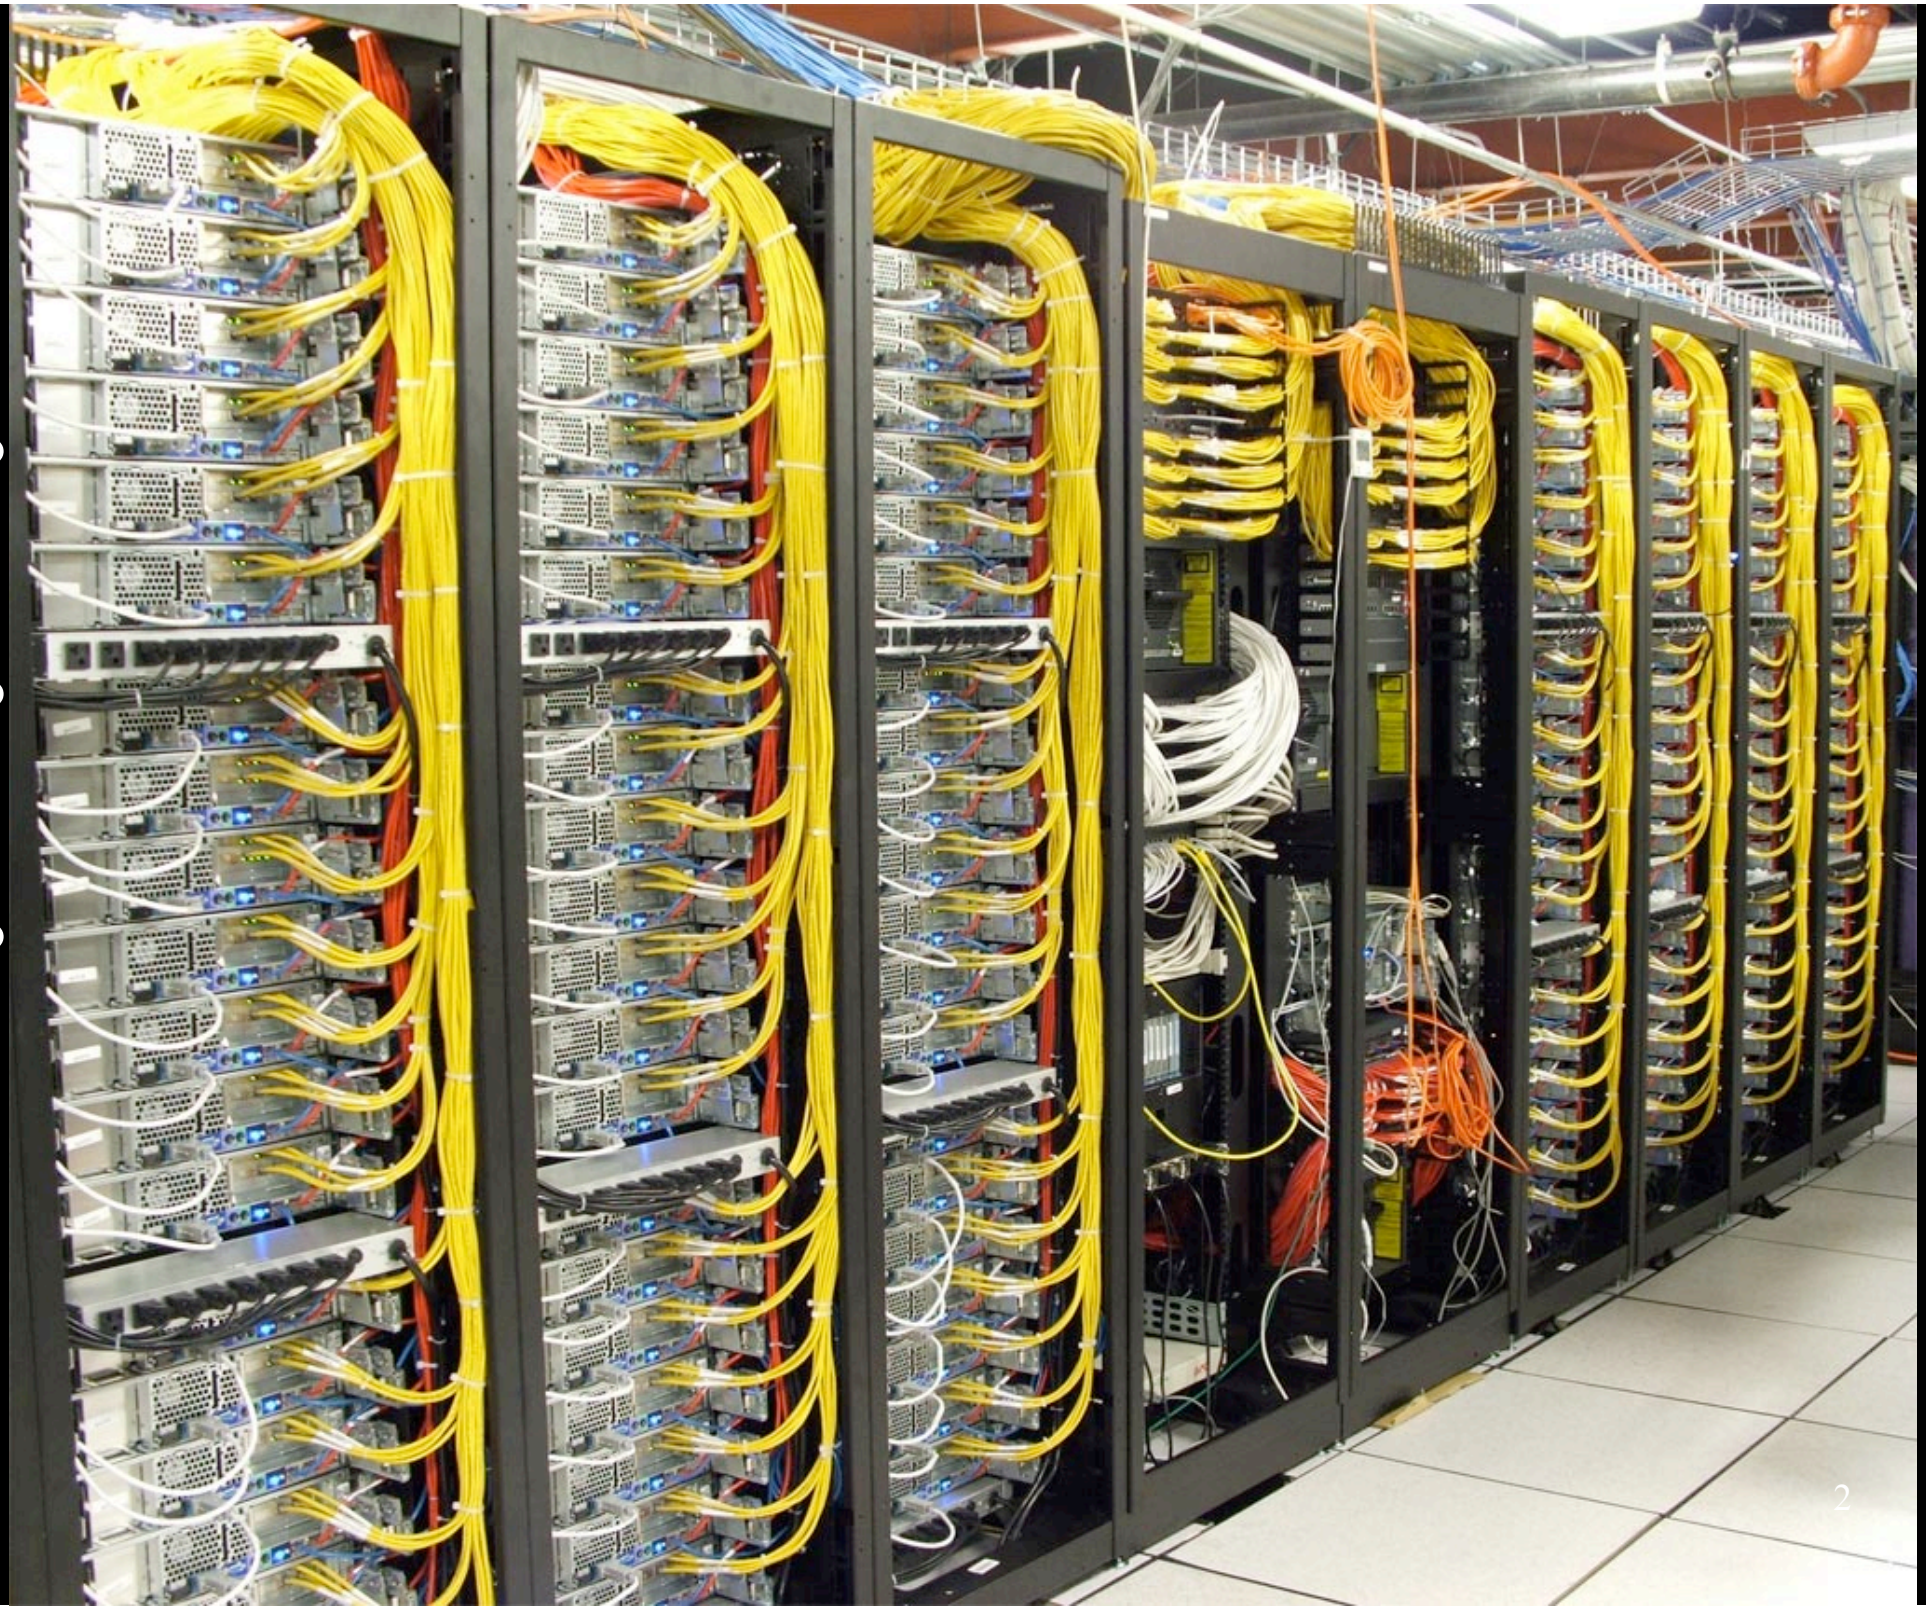

Emulab Nodes

- Physical nodes
- Users have root
- Space/time shared

Artifacts from previous experiment may persist on node

Node Corruption

Why Reset State?

- Experiment fidelity depends on starting fresh
- Unacceptable for security sensitive experiments
- At the very least, artifacts from previous experiments are irritating

Emulab's Current Method

- Control server forces reboot and directs node re-imaging over network
- Network is shared with other nodes

State reset is not guaranteed and is not tamper-proof

Goals

- Must reset node state during other active experiments and regardless of what state the node is left in
- Must be flexible for many boot paths
- Must scale to size of testbed

Solution: Trusted Disk Loading System (TDLS)

If the experiment is created successfully, node state is reset

Contributions

- Cryptographically verifiable method of resetting physical node state
- Flexible and secure reloading software scalable to size of testbed

Node Reloading

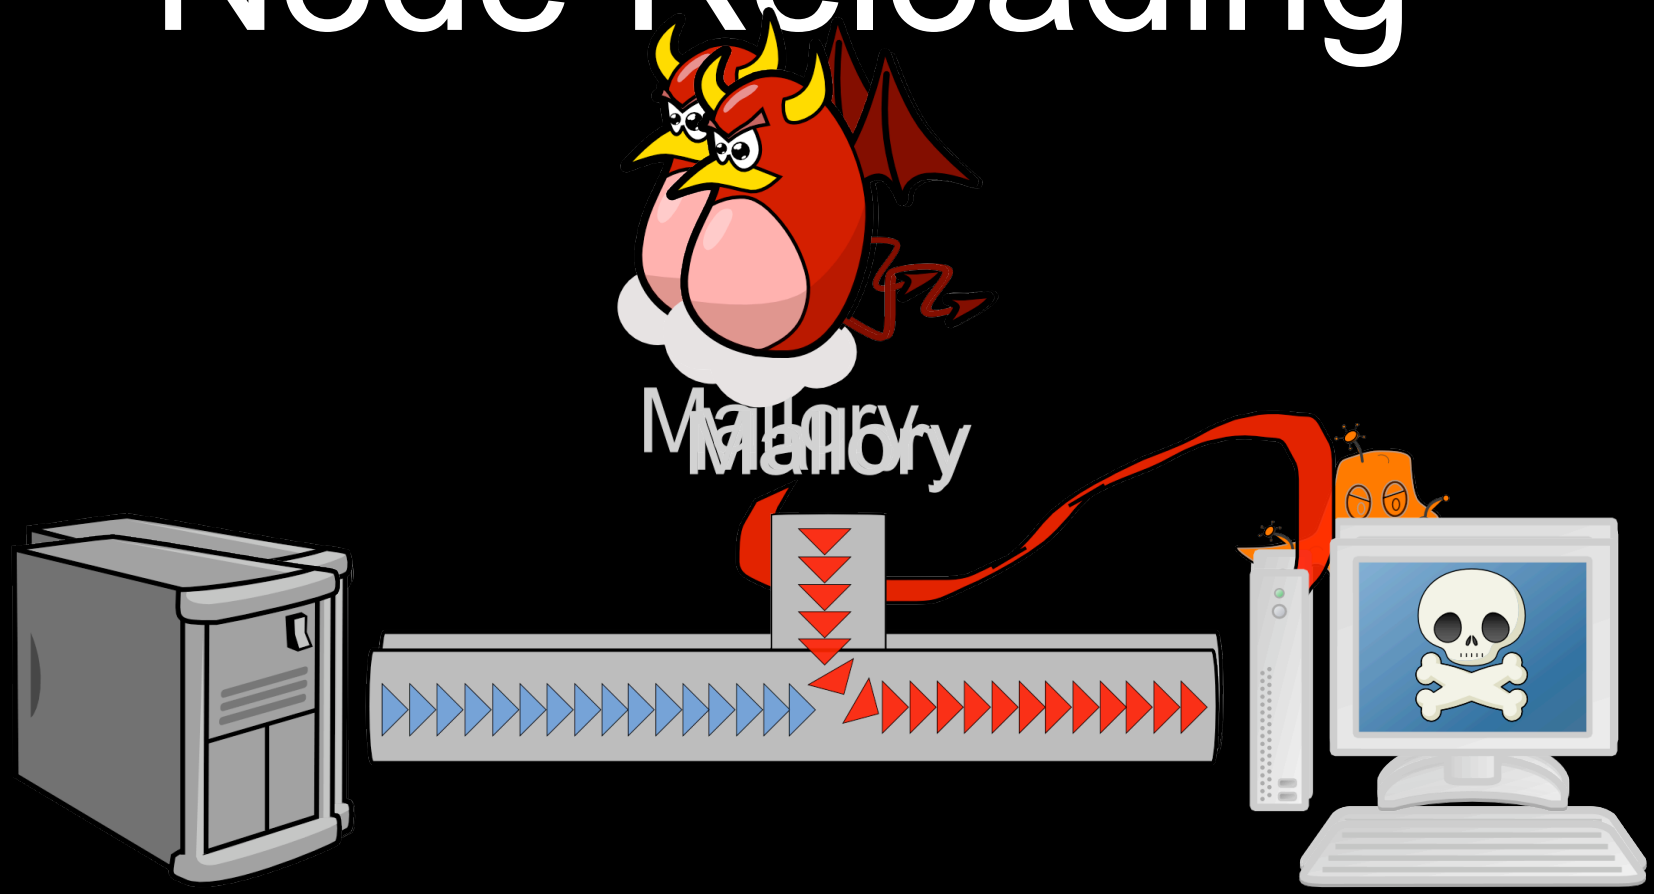

download

boot ROM

TDLS Fundamentals

- Establish trust
- Verify every stage of node reloading with control server

The Trusted Platform Module is the perfect tool for such objectives

Trusted Platform Module (TPM)

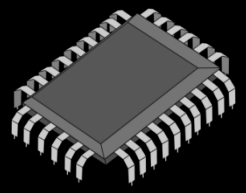

- Secure key storage
- Measurement
- Remote attestation (quotes)

Secure Key Storage

- Keys are always encrypted before they leave the TPM
- Keys are only useable on the same TPM with which they were created
- Control server can identify nodes by the public portion of these keys

TDLS Fundamentals

- ✓Establish trust
- Verify every stage of node reloading with control server

Trusted Platform Module (TPM)

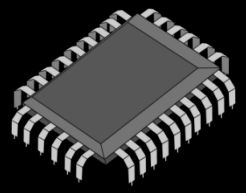

- Secure key storage
- Measurement
- Remote attestation (quotes)

Measurement

- Platform Configuration Registers (PCR)
  - TPMs generally have 24 PCRs
  - Holds a hash
  - PCRs can only be modified through extension
  - Extending:

    PCR = hash(previous value of PCR + a new hash)

- Measuring is when we hash a region of memory and extend a certain PCR with the resulting hash

Secure Boot Chain with TPM

1. Immutable part of BIOS measures the rest of BIOS
2. BIOS measures boot device
3. Boot device then measures whatever it loads
4. etc.

Remote Attestation (Quotes)

- TPM packages up the desired PCRs and signs them
- Tamper-proof as it is signed by the TPM
- Very easy to differentiate between a genuine quote and arbitrary data signed by TPM

TDLS Fundamentals

- ✓Establish trust
- ✓Verify every stage of node reloading with control server

TDLS Reloading

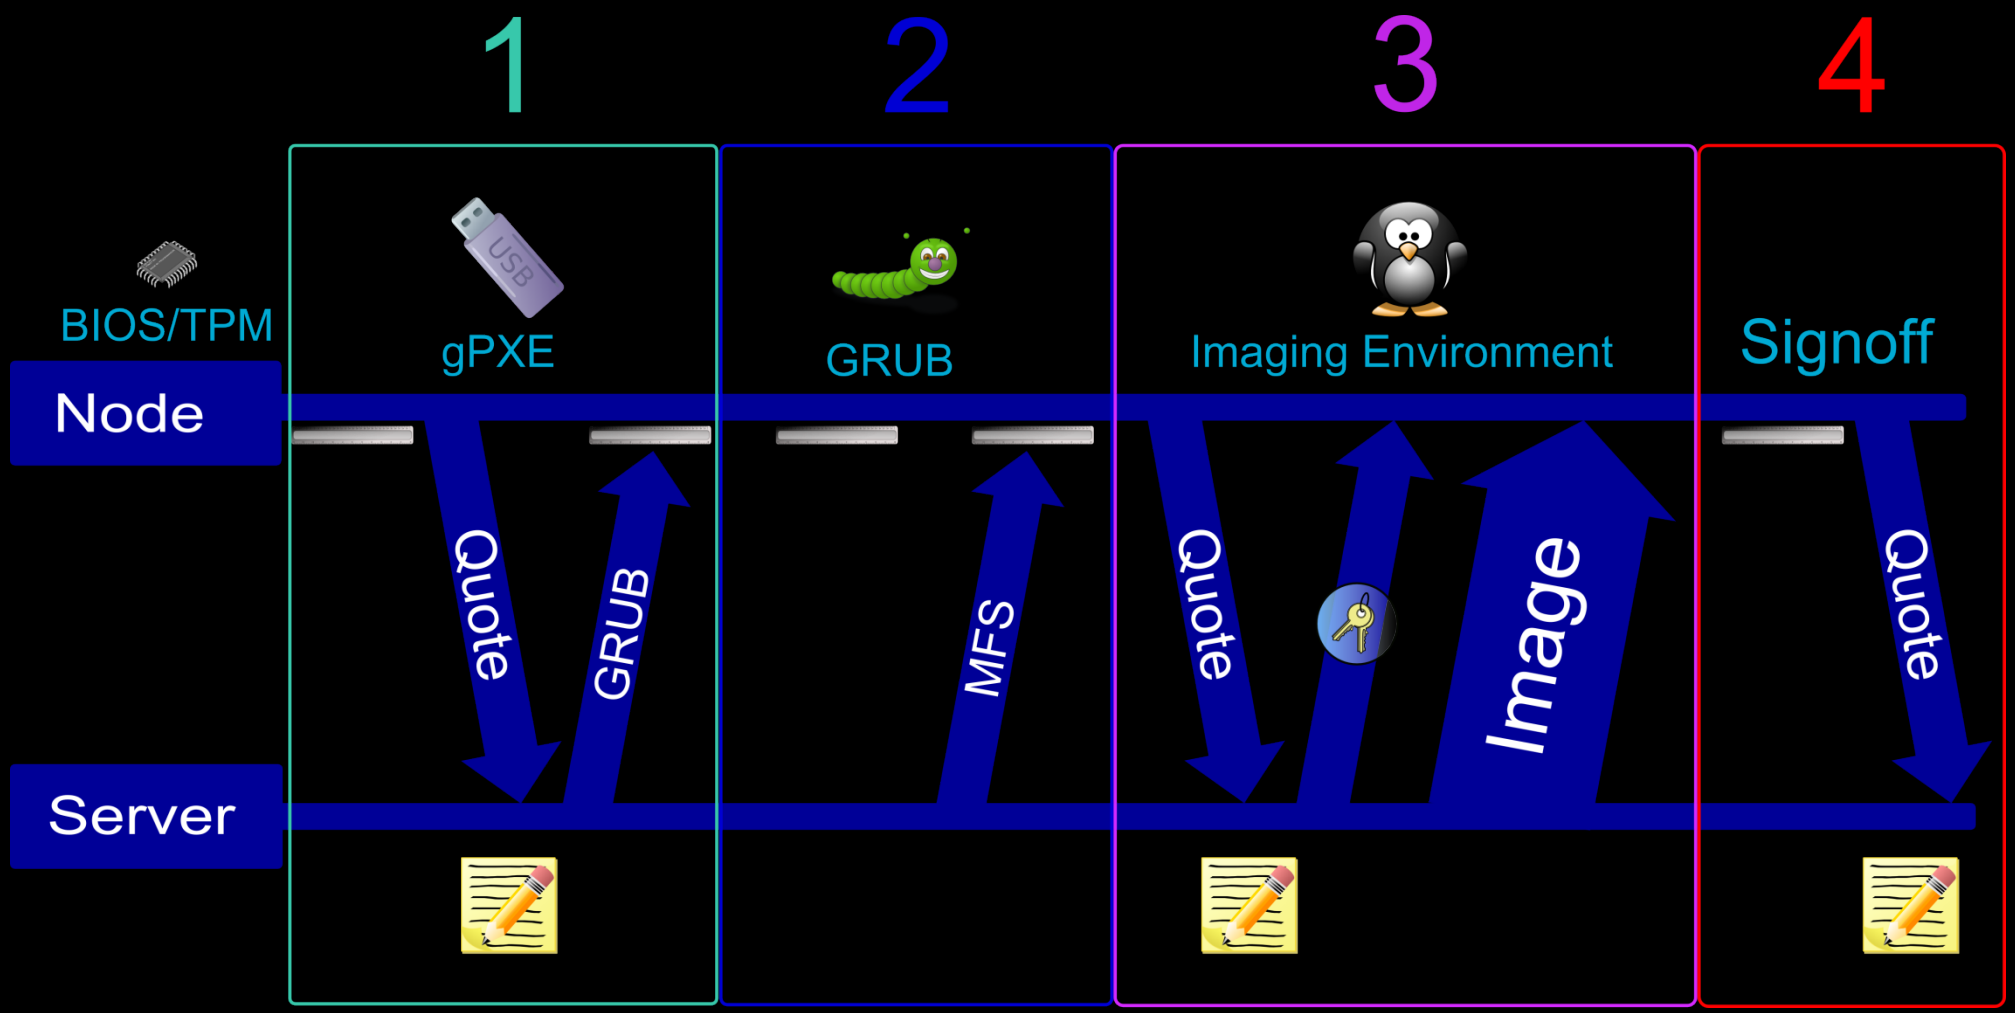

Starting the chain: Booting to PXE

- PXE ROMs aren’t TPM aware
- PXE ROMs won't check-in with the control server

Boot to USB dongle with gPXE

Stage 1: gPXE

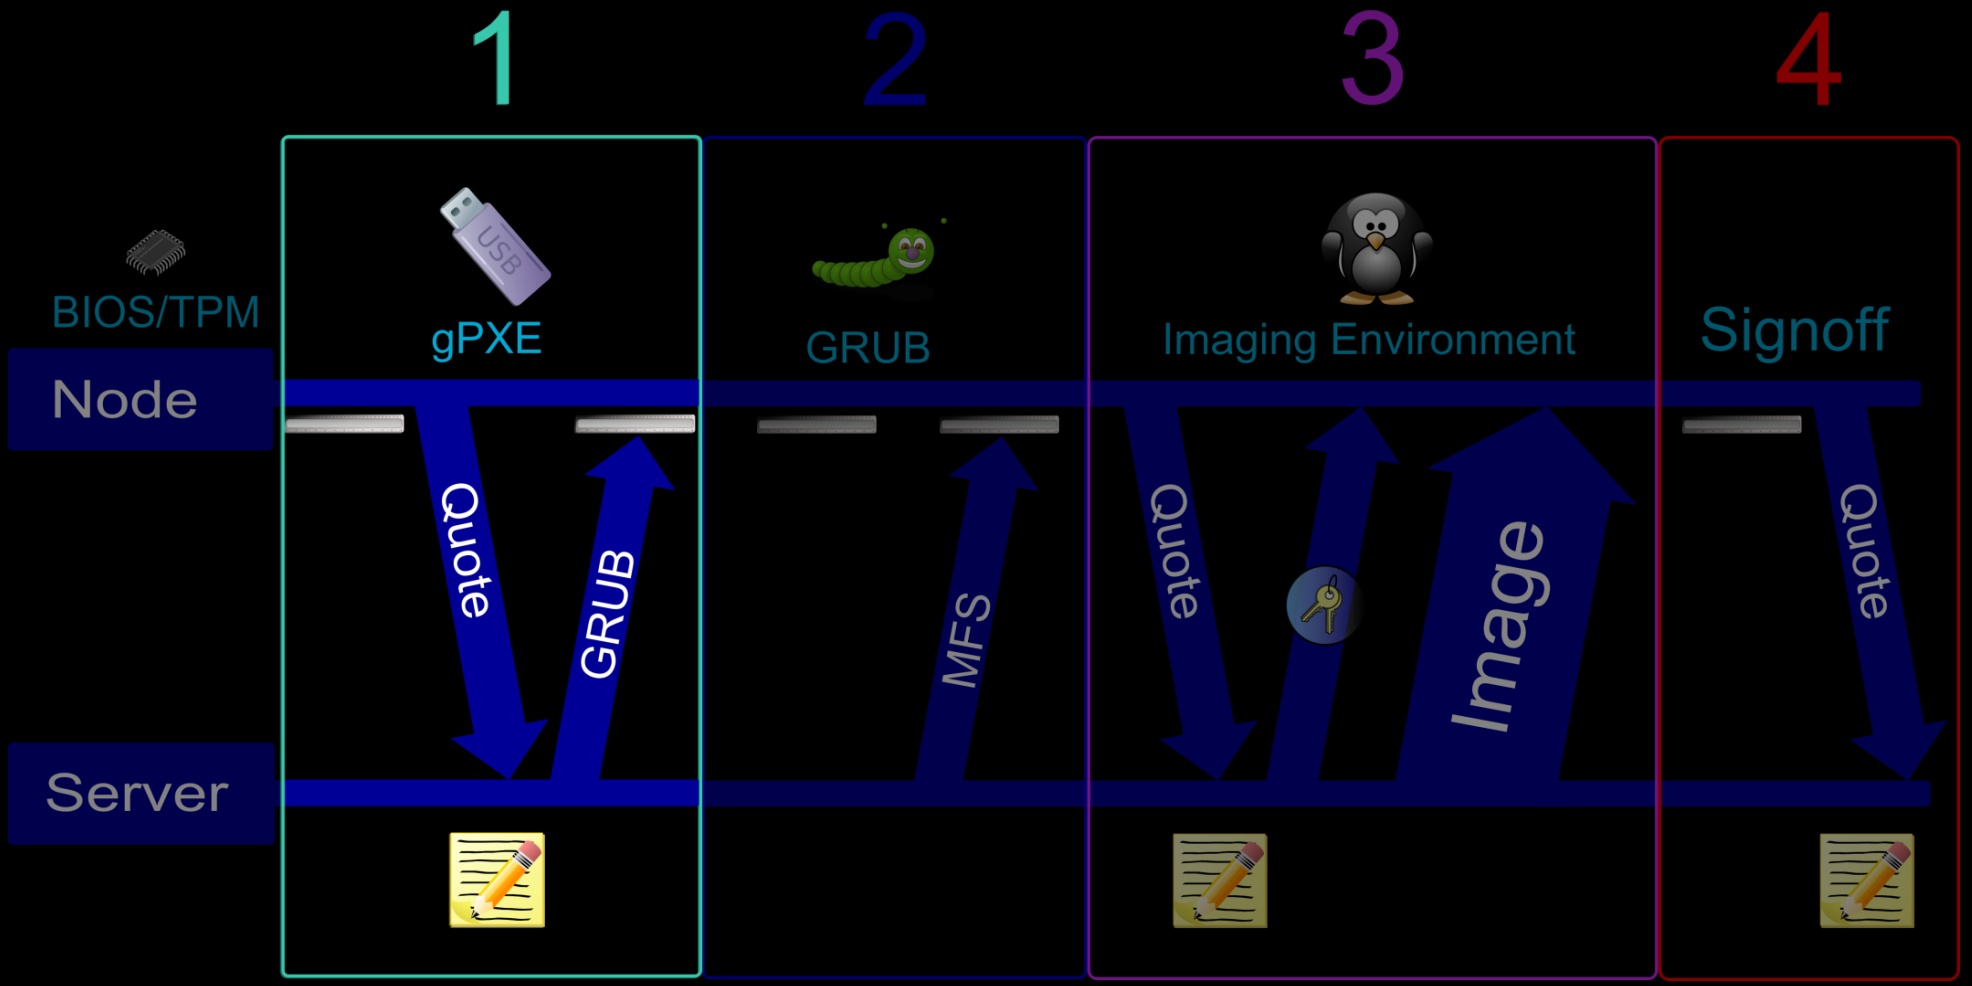

- Measured by BIOS
- Embedded certificate authority for server authentication
- Sends a quote to control server

Checking Quotes

- Different stages are measured into different PCRs
- Quotes contain a nonce from the server to guarantee freshness
- The TPM signature over the quotes are verified
- Server compares every PCR in the quote with known values in the database

Incorrect Quotes

- An incorrect PCR means something was modified
- Failure to send a quote before a timeout is treated as a verification failure
- Control server cuts power to the node and quarantines it

Stage 2: GRUB

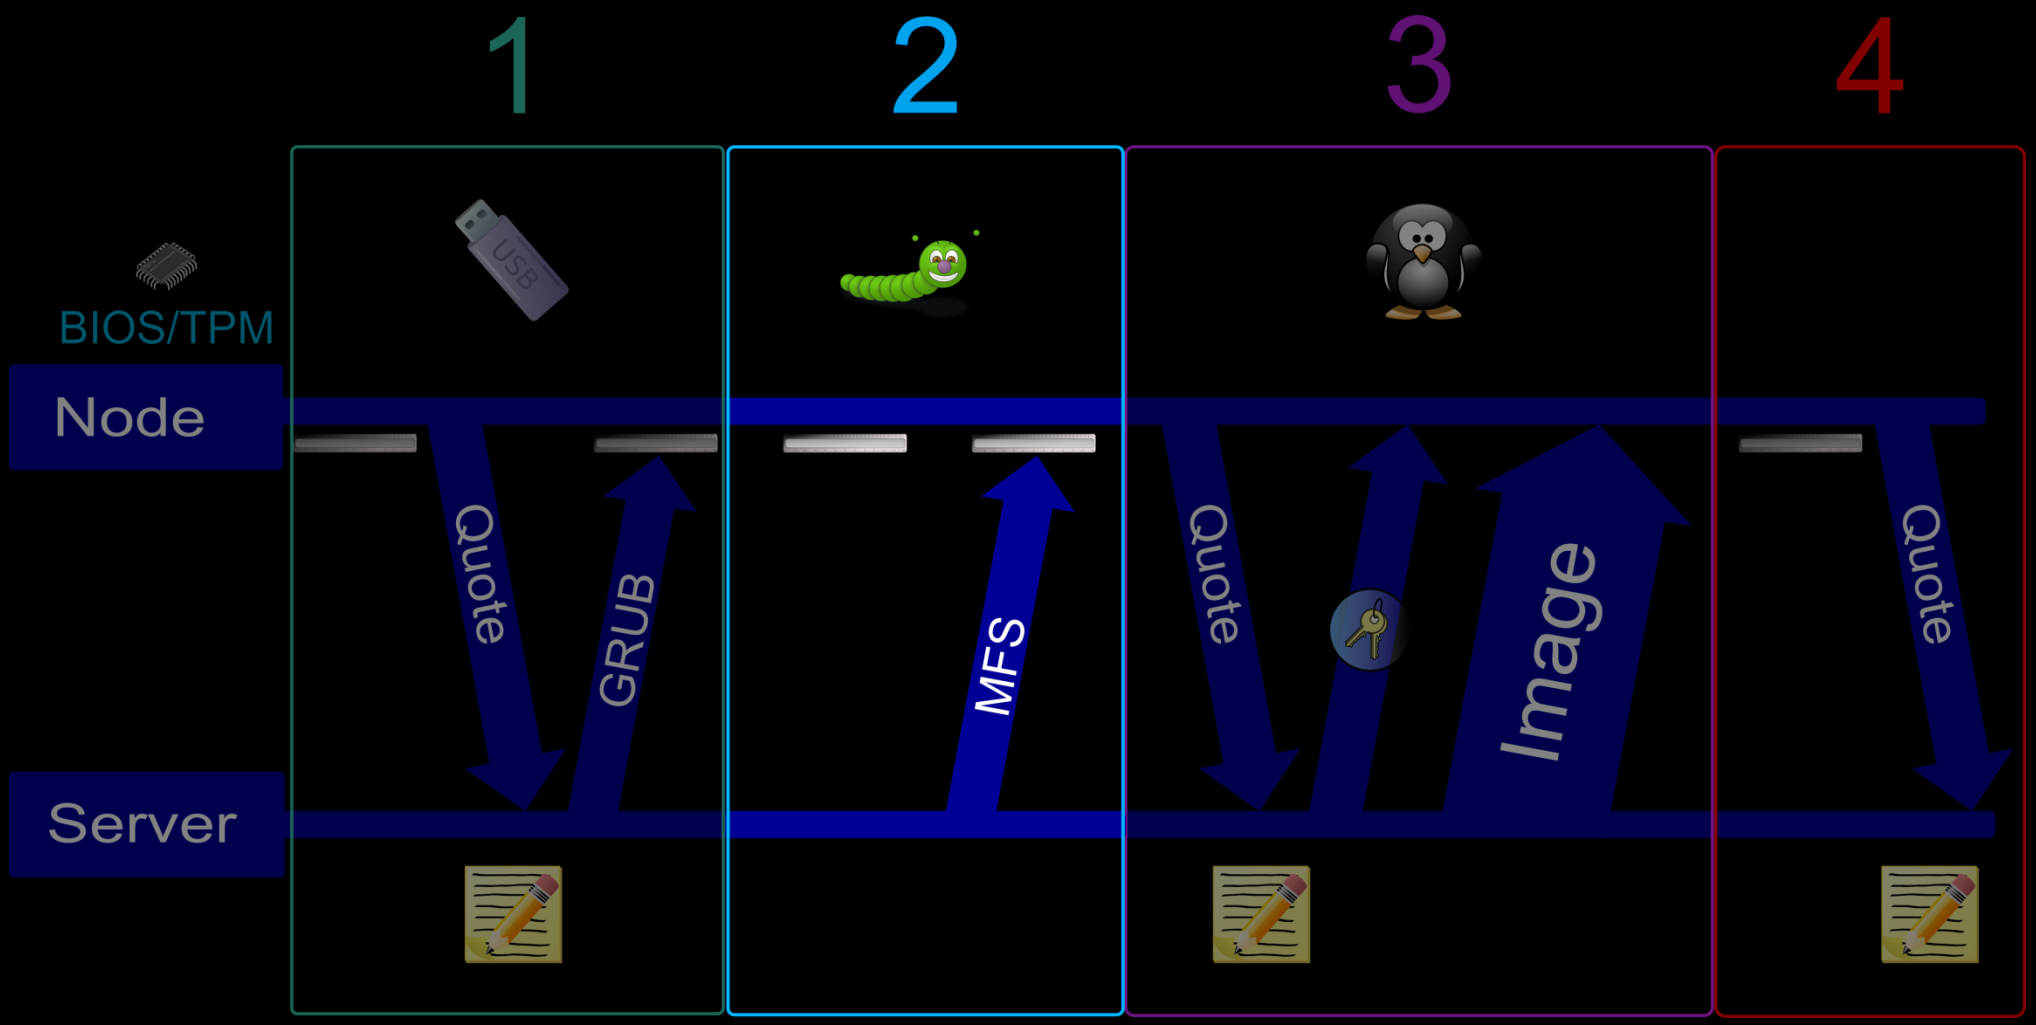

- Retrieves, measures, and boots the imaging MFS
- Will boot to disk when necessary

Sensitive Resources

- Control server closes monitors a node's progress via quotes
- A node can only receive sensitive resources (decryption keys) in a particular state

Stage 3: Imaging MFS

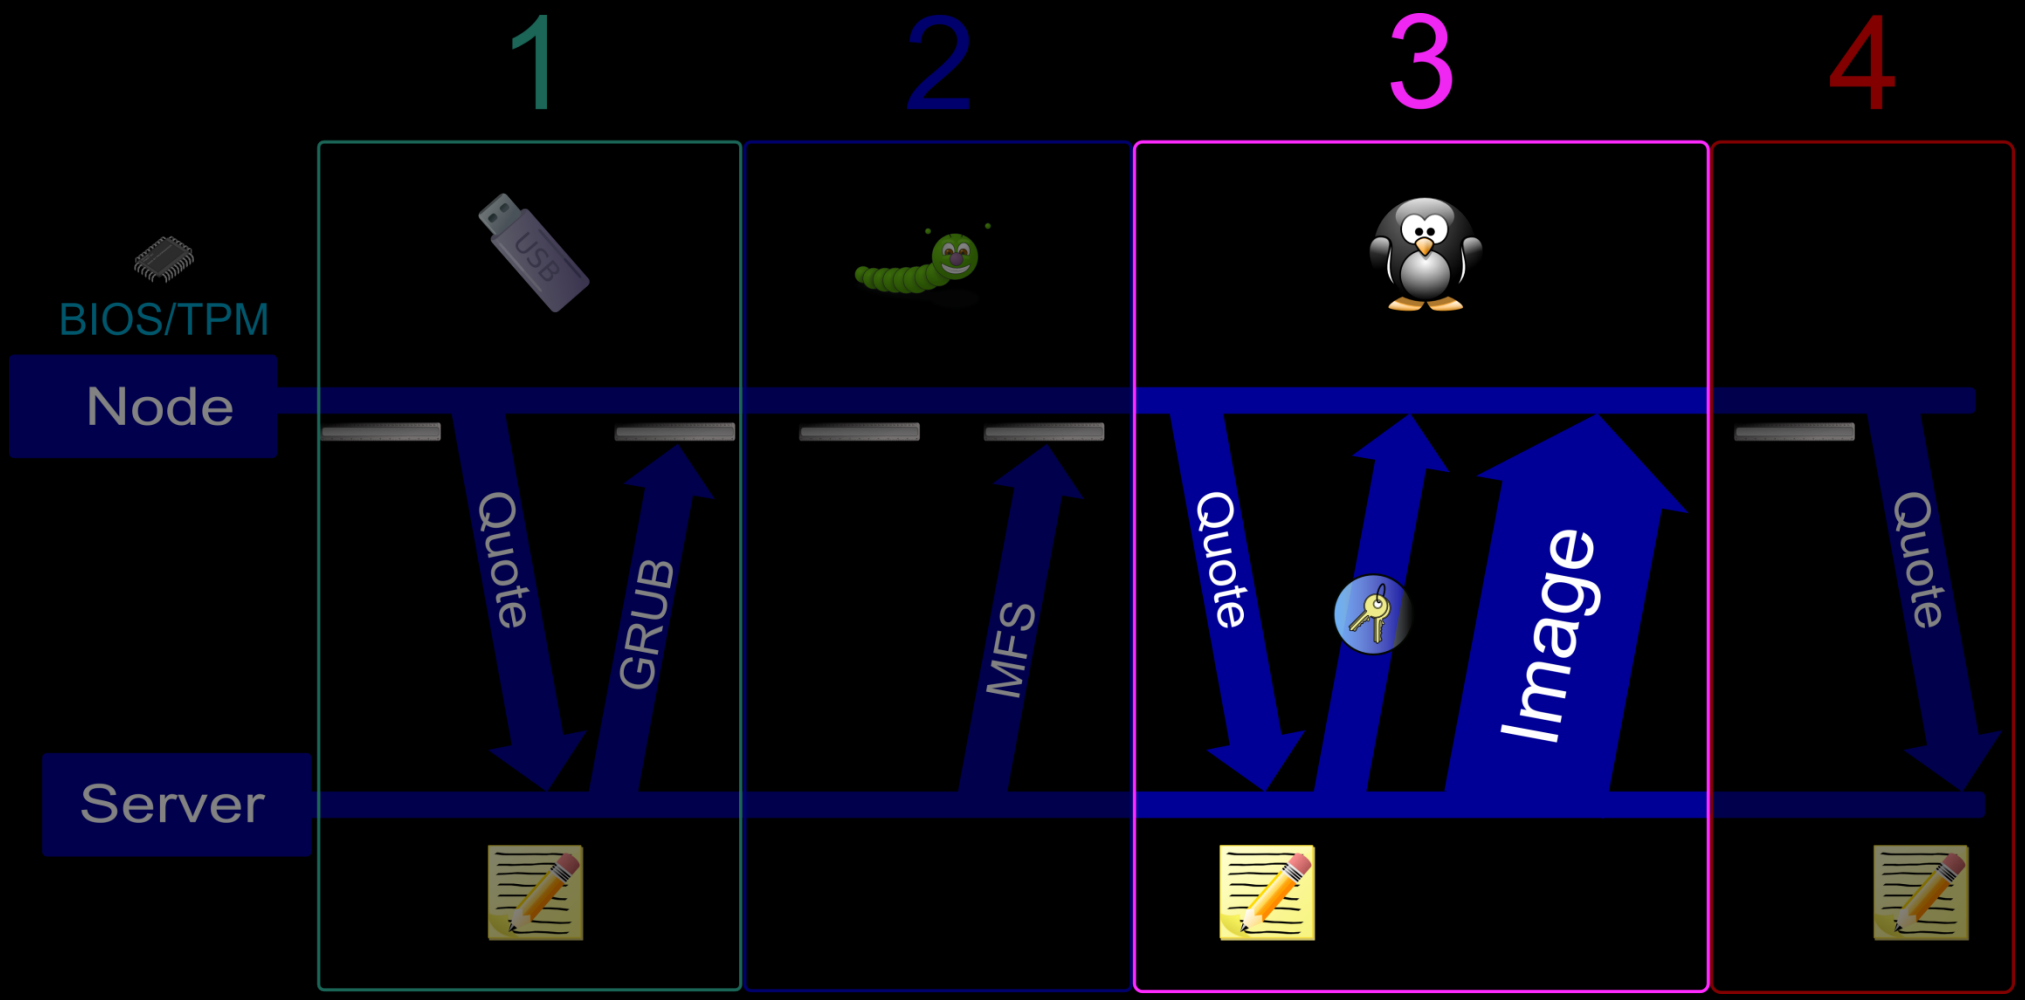

- Sends quote covering everything
- Writes the encrypted image to disk

Stage 4: Signoff

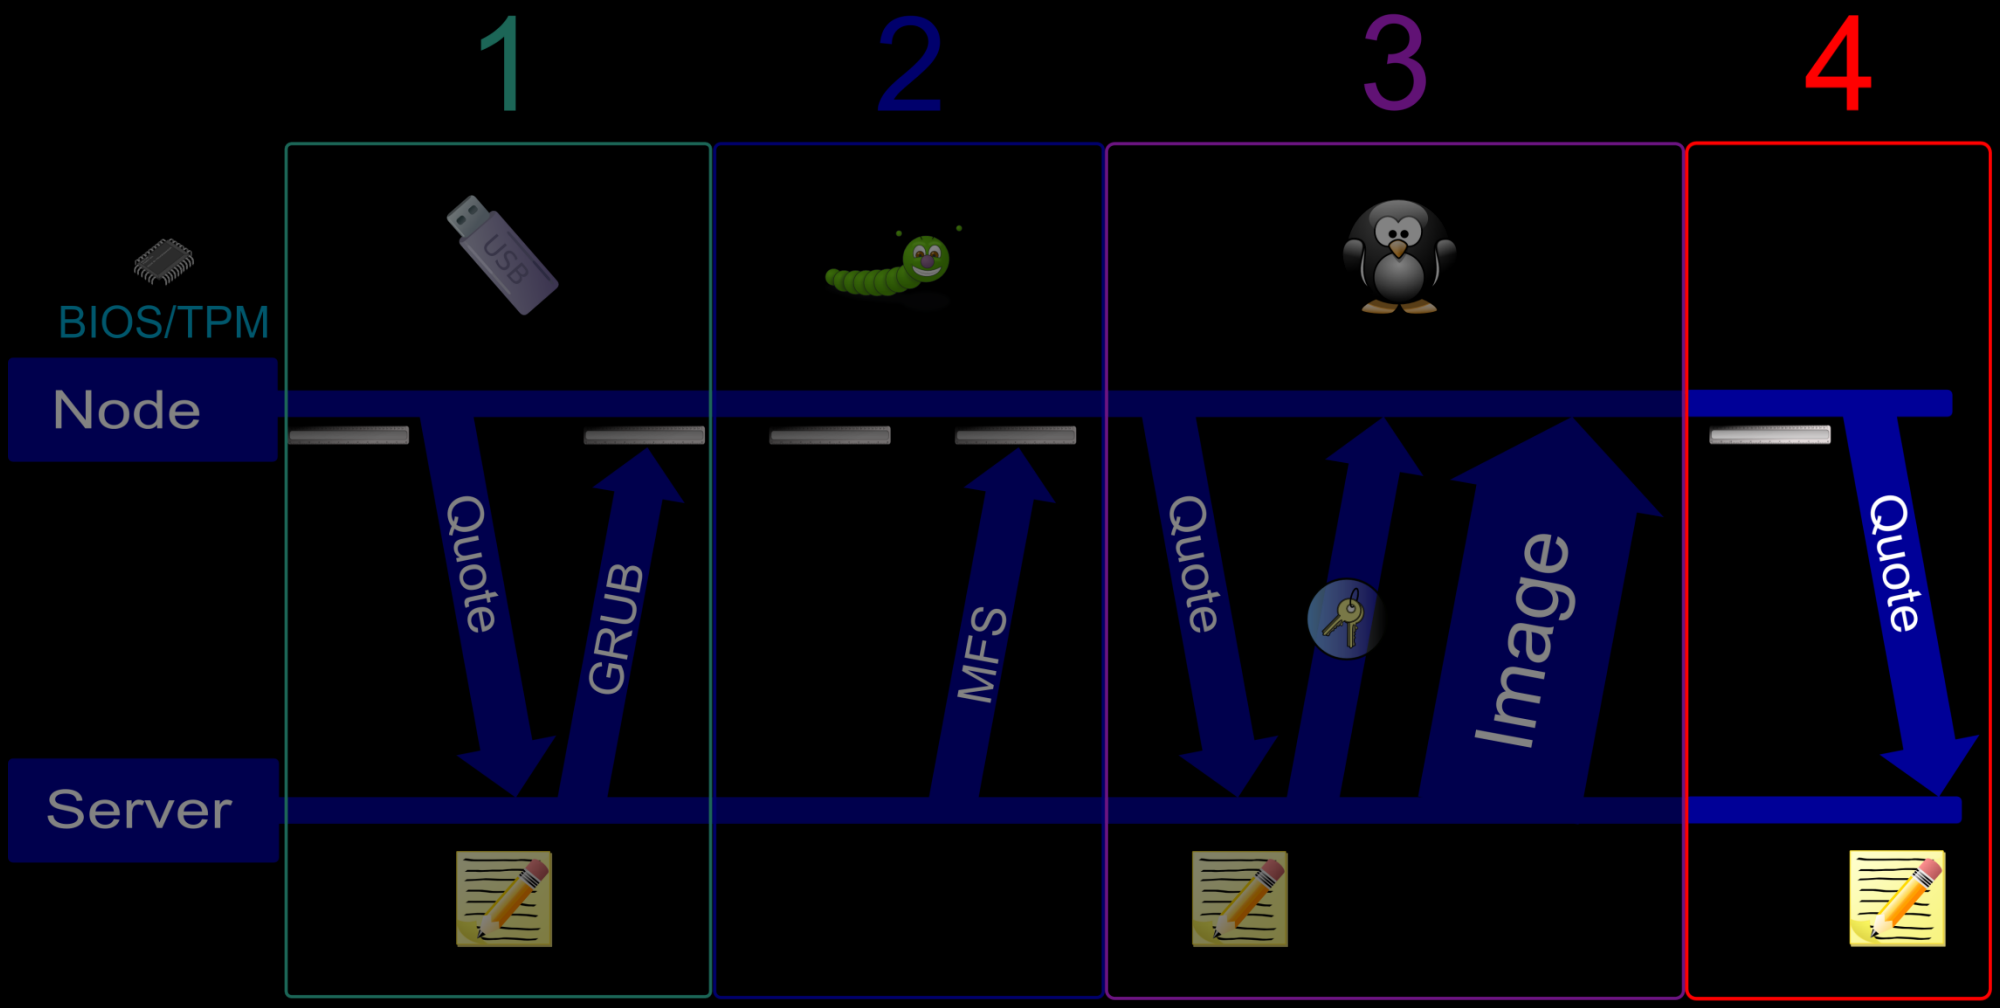

- Disk is imaged
- Extends known value into designated reboot PCR
- Marks the end of the trusted chain

Attacks That Will Fail

- Any boot stage corruption
- BIOS code or configuration modifications
- Injecting new stages

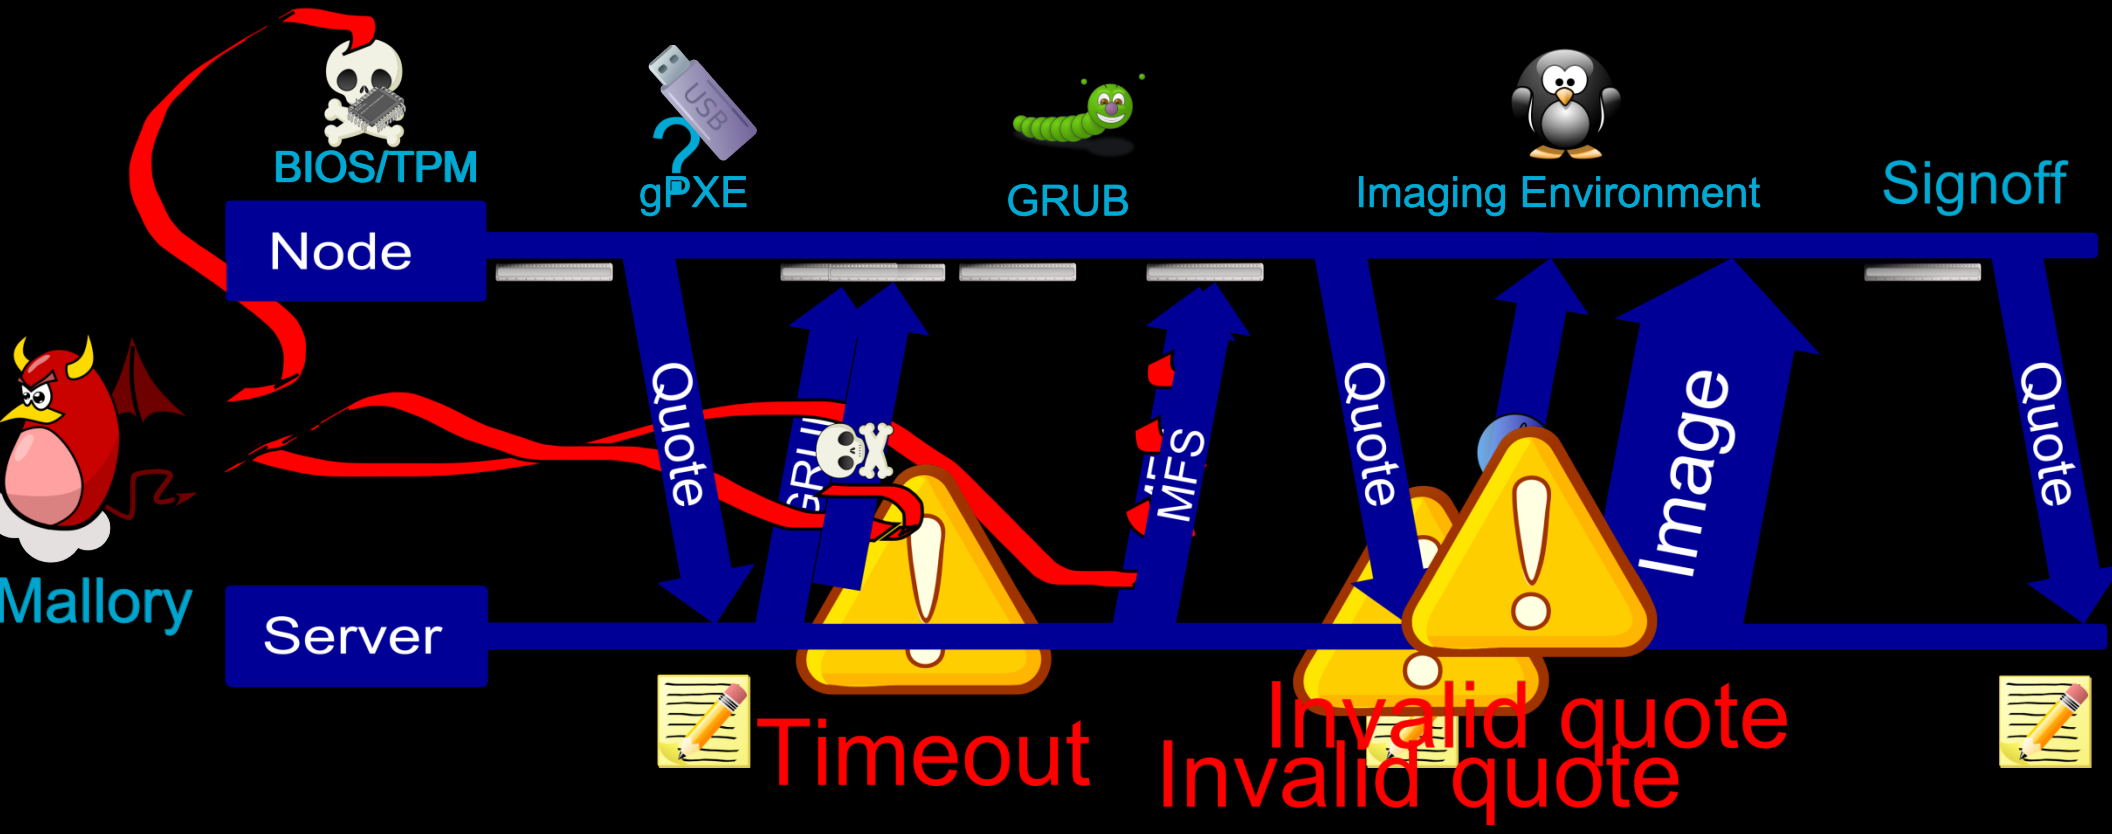

What this means

We win

Summary

- Node state must be fully reset in a secure way
  - Some testbed properties make this very difficult
- Using the Trusted Platform Module
  - Establish trust between the node and server
  - Verify every stage of bootchain
- Trusted Disk Loading System
  - Tracks node progress with quotes
  - Guarantees node state is reset
- If any check fails, the experiment creation will fail

Future Work

- Enable experimenters to verify node state
- Refine the violation model
- Integrate with Emulab UI
- Deploy on 160 TPM-enabled nodes at Utah

Questions?

ccutler@cs.utah.edu

emulab

http://www.emulab.net

Solution: Trusted Disk Loading System

- If the experiment is created successfully, disk is imaged as expected
- Scalable to size of testbed
- Flexibility for the addition of many boot-paths
- Prototype

Guarantees

- If any check fails, the experiment creation will fail
- Disk is imaged as specified