DISA

Defense Information Systems Agency

A Combat Support Agency

Speed

Richard Hale
Defense Information Systems Agency
Chief Information Assurance Executive
October 28, 2009

1. What's My Job (and possibly your job) and Why Do I Think We're Here?

2. Bad Guys

3. Handling Bad Guys
   - Keep them out in the first place (A)
   - Detecting, diagnosing, reacting to cyber attack (in case "A" fails)
   - Keep them out in the first place (B)

4. Bad Guys Again

It's All About Mission

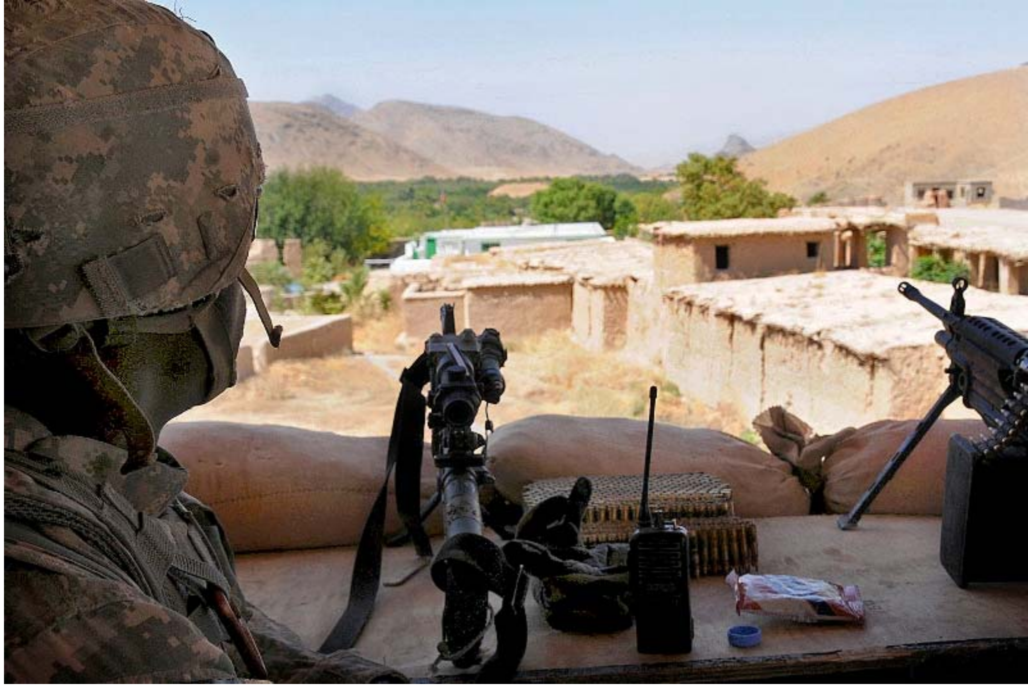

Forward Operating Base Mizan, Afghanistan, Sept. 10, 2009

Vehicle Patrol Base Badel, Konar province, Afghanistan, May 9, 2009

Rigged alternate method zodiac or RAMZ jump, Operation Southern Partner, Caribbean Sea, June 9, 2009

Corollary: “Security” Isn’t the Point

Information Assurance

My Job Part 1:

Ensuring that DoD personnel *and DoD's mission partners* can depend on information and on the information infrastructure in the face of cyber warfare by a capable adversary

Job 1 restated.

We need missions to be *dependable* in realistic operational environments (in the face of cyber warfare)

(Aka Mission Assurance)

Or

Mission Dependability

(in the face of cyber warfare)

My Job 3:

Ensuring that DoD and its mission partners can keep a secret (when we/they want to)

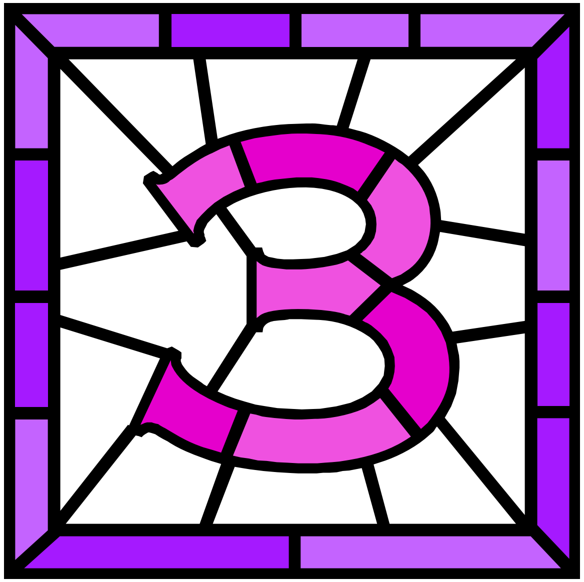

My Job 2:

Doing jobs 1 & 3 while enabling (safe) sharing with the broadest set of partners and possible partners

1. Dependability in the Face of Cyber Warfare

3. Keeping a Secret

2. Safe Sharing

Bad Guys

Kinds of Bad Guys I Worry About

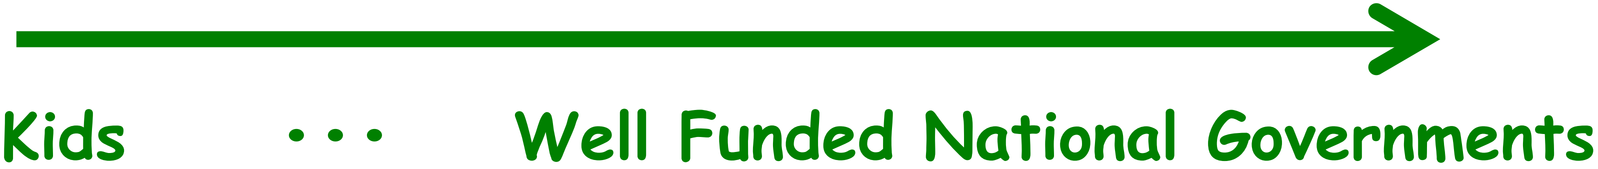

Kids . . . Well Funded National Governments

(Every Kind)

Threat (and/or Technology) Realities

- Attacks are ***easy, cheap, scalable, & can be developed fast***

(Bad Guys are fast)

- Bad Guys don't need much infrastructure so they don't get wedded to it
  - *(Since attacks are easy, cheap, scalable...)*

(Bad Guys are agile)

Keeping Bad Guys Out (1): *Reduce Attack Surface*

Most Basic Goal: *Configure Every Computer Securely, Keep It Configured Securely As Things Change, and Ensure the Right People Know This is So (or Not So)*

Achieving That Most Basic Goal

(Stuff You Already Know)

- ***Very hard*** to configure properly manually
- ***Impossible*** if one manages many computers
- ***Very hard*** even to find every computer one owns
- ***Very hard*** to figure out manually in what configuration one finds any particular computer
- ***Very hard*** to manually report configuration
- ***Very hard*** to change configuration manually fast in response to new vulnerabilities

Conclusion: Solution to Most Basic Problem is . . . *Automation*

(You might have seen this coming)

- Configuring
- Changing configuration
- Measuring configuration from within a device
- Measuring configuration externally as a double check
  - (Remember the bad guys . . . compromised boxes sometimes don't tell the truth about themselves)
- Reporting configuration continuously

What DoD Needs Here (1)

- As a developer of secure configurations and perhaps of methods of measuring configuration:
  - Need to define it once, then have anyone, with whatever configuration tool, consume it
    - Interoperable configuration content
  - Also need to know the configuration is secure, (by knowing how it maps to the definition of "secure"), and so might others who weren't involved in inventing the configuration
    - Standard for deriving configuration from security controls (more later)

What DoD Needs Here (2)

- As a purchaser of tools:
  - Must be able to consume the standard content, push a button and have all my machines configured properly
  - Need to understand what policies the content is deploying
  - Want good configuration tool competition
- As someone who wants to understand the state of configuration of my computer, those in my organization, those in my subordinate organizations, those in my partner organizations:
  - Must consume & understand the most current measurement content
  - Consume and/or report the most current measurements in a way anyone who needs to can understand

What DoD Needs Here (3): Unique, Standards-Based Device Identity

- Must be able to identify a particular device
- Must be able *correlate* reports about this device from configuration tools, multiple measurement tools, and other tools
- Unique, standards-based device identity used by every tool that configures, measures, reports, etc., is critical

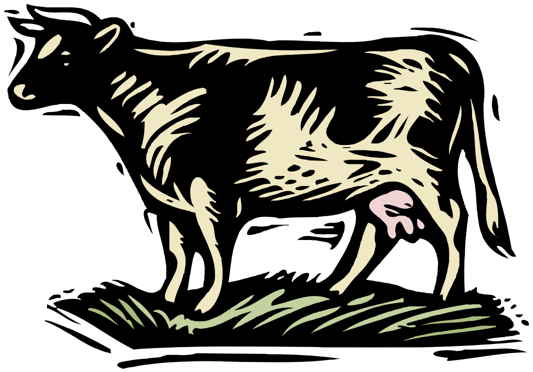

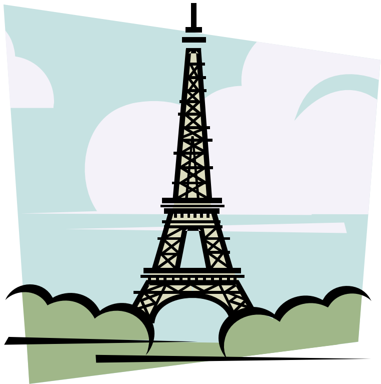

Standards-based, Unique Device Identity Isn't Just A Security Automation Problem

- It's critical across all other information technology management approaches
- Policy-based network management, application management, security automation, etc.

- (These are all in the service of dependable service for customers)

Reduce Attack Surface (2): Shield Vulnerabilities that Can't Be Configured Away

Data Centers Are Complicated Places:

Is Your Application Really Behind/Within the Perimeter Policy You Need?

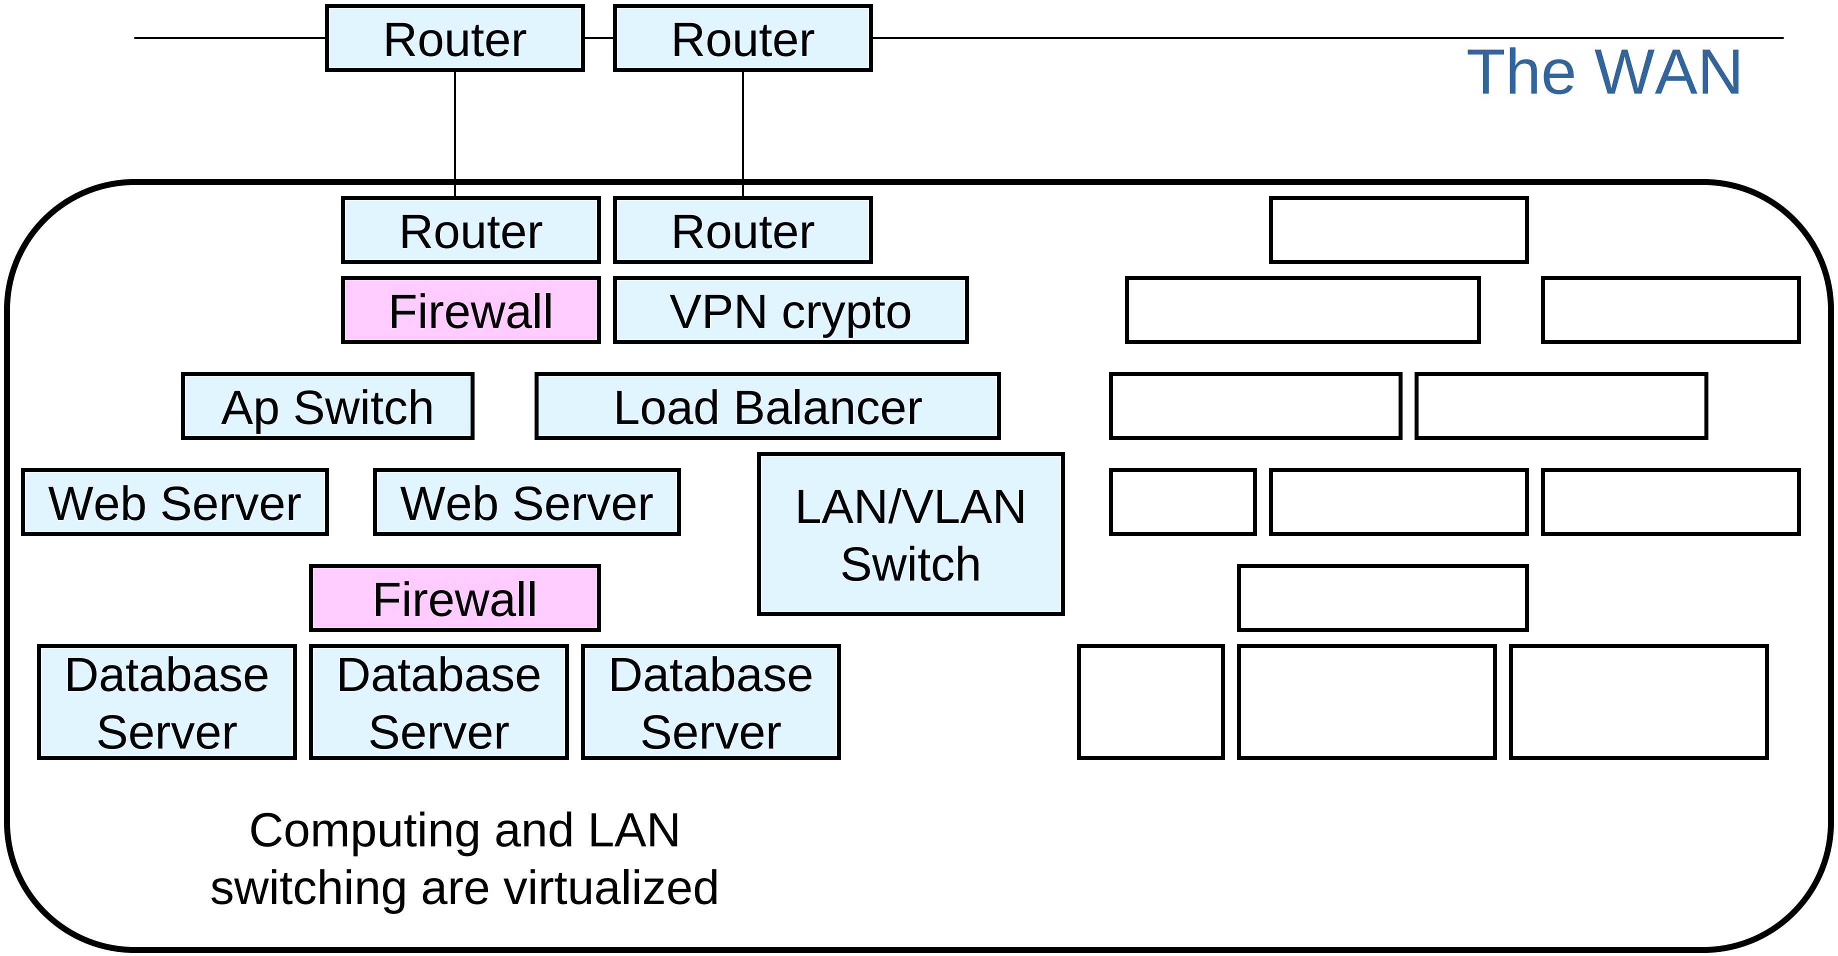

The perimeter security policy in a data center will also depend on things like switch configurations/VLAN configurations, IPsec VPN configurations, load balancer configurations, and the like

This Whole Mass of Stuff in the Data Center Must Be Configured Securely, Changed, Measured, Reported

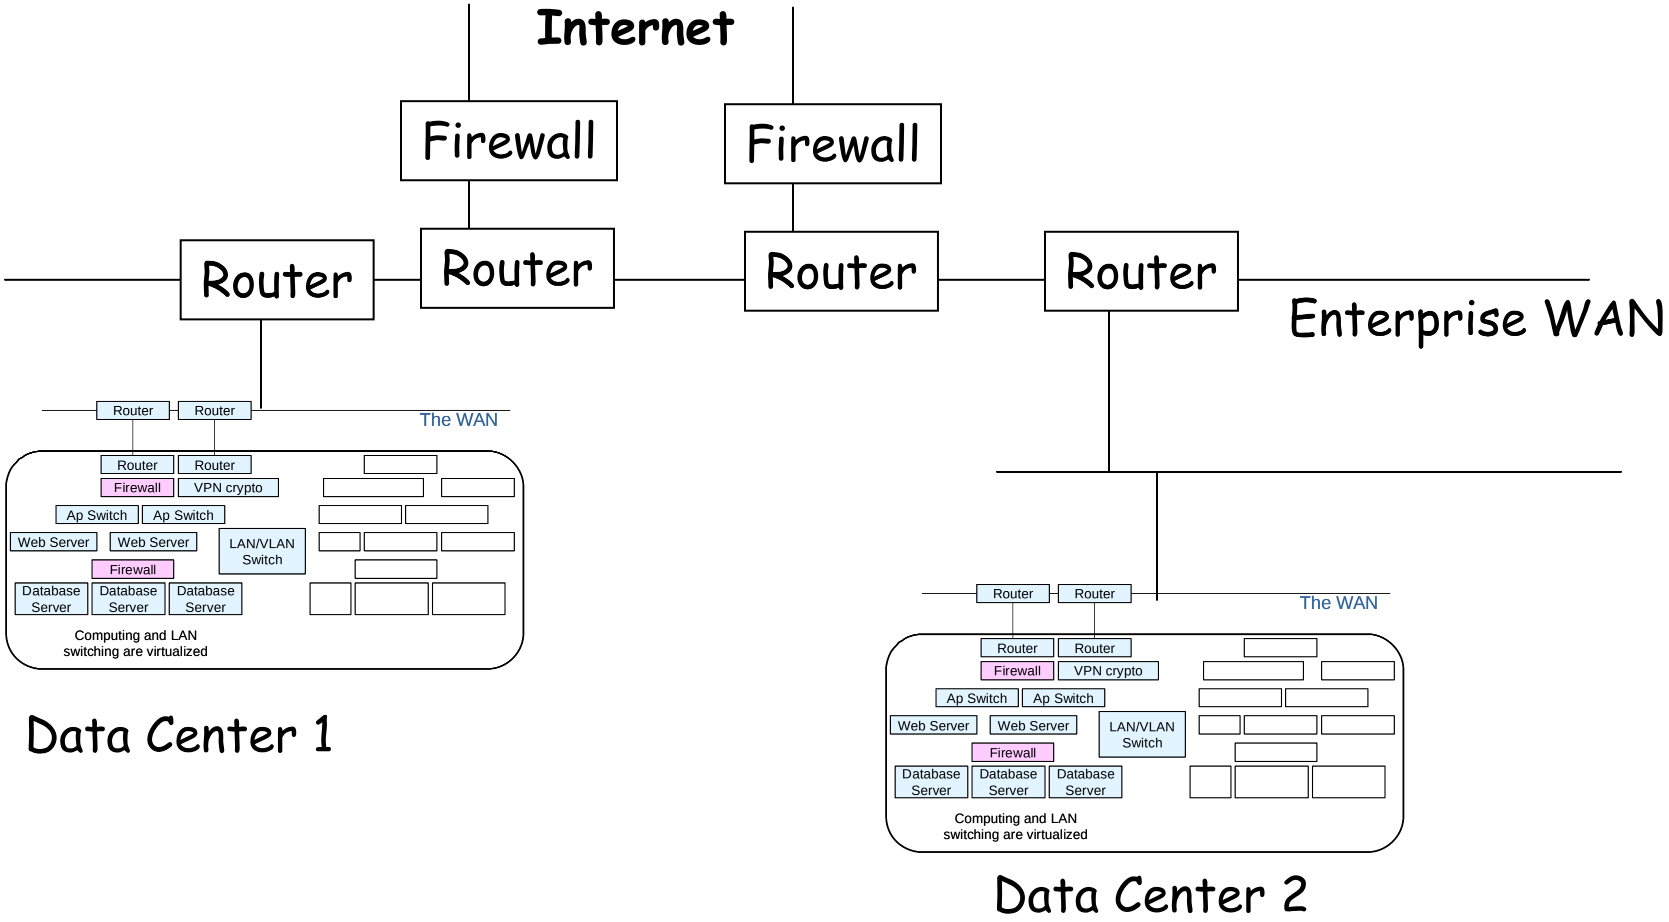

Internet

Firewall

Firewall

Router

Router

Router

Router

Enterprise WAN

Router

Router

The WAN

Router

Router

Firewall

VPN crypto

Ap Switch

Ap Switch

Web Server

Web Server

LAN/VLAN Switch

Firewall

Database Server

Database Server

Database Server

Computing and LAN switching are virtualized

Data Center 1

Router

Router

The WAN

Router

Router

Firewall

VPN crypto

Ap Switch

Ap Switch

Web Server

Web Server

LAN/VLAN Switch

Firewall

Database Server

Database Server

Database Server

Computing and LAN switching are virtualized

Data Center 2

So?

Attack Surface Reduction Automation: *What We Really Need (almost)*

- **Help a human** define policy (in English) for a globally-load balanced, multi-tier application, *parts of which inherit security control conformance from stuff others control in the data center,* on the enterprise WAN, and at the WAN/Internet boundary, partly based on government standard security controls. Properly account for topology/relationships. **Automate** what-ifs in defining the policy
  - **Help a human** understand the policy
- **Automatically** configure everything (physical computer, virtual computers, application, data center, WAN, WAN perimeter, etc.) to this policy
- **Automatically** measure the policy that's really on the ground
- **Automatically** understand the as-built policy, and **help a human** understand it too
- **Automatically** figure out where/why it doesn't match the one we set, and **automatically** propose configuration/policy changes
- **Automatically** deploy these, then remeasure, etc.

Internet

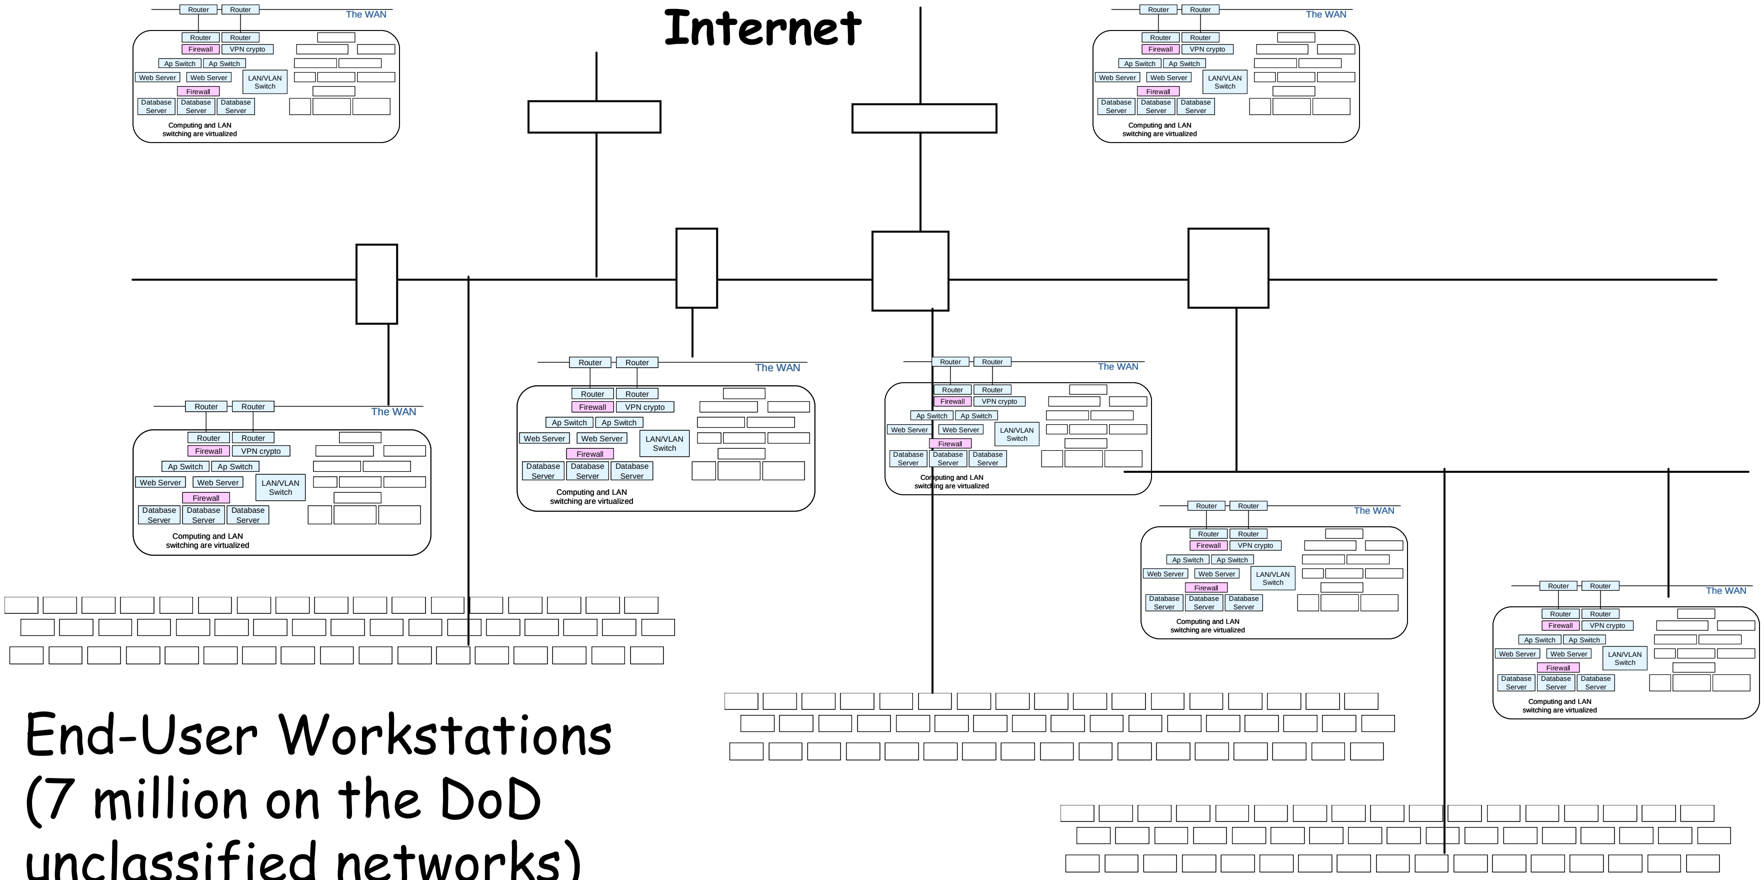

End-User Workstations
(7 million on the DoD
unclassified networks)

Somewhere in Here Are Poorly Understood Transitive Trust Relationships that Span firewalls/VPNs, etc.

Attack Surface Reduction Automation: **Measure, Correlate, and Find All Those Transitive Trust Problems**

- Help develop policy change recommendations for each organization that “owns” part of the problem
- Deploy these so organizations can automatically apply, measure, report
- Or at least help someone be aware of policy differences between organizations that may matter

Now We Have…

(Consistent) Machine-to-machine attack surface reduction

Policy-based attack surface reduction

(lots of other policy-based management things)

Security Automation and Sharing

DoD Mission Reality:
Everything We Do Is a Coalition

(you get the idea)

Joint

Coalition, Joint

Interagency, Coalition, Joint

Industry, Interagency, Coalition, Joint

State & Local, Industry, Interagency, Coalition, Joint

NGO, State & Local, Industry, Interagency, Coalition, Joint

State & Local of Other Countries, NGO, State & Local, Industry, Interagency, Coalition, Joint

Establishing Conditions for Partnership: Share Configuration Measurements

- To enable sharing, partners need to trust each other…and trust that each partner isn't introducing undue risk into the partnership.
- **Want a Consistent Policy Across All Mission Partners** (and know this is so)
- **Or, Want understanding of partner's policy and policy conformance so (automated?) risk decisions about how much to share, how much to open certain defenses…**
- Then want to know when something important changes so my policy can shift (automatically?) to contain damage or whatever.
- Information Interoperability is critical

At This Point We've Really Reduced Our Attack Surface, and Are Keeping Reduced

Capable Bad Guys Are Still Clever, Persistent, and Not Wedded to Particular Approaches (remember *fast*, scalable, etc?)

Cyber Attack Detection, Diagnosis, Course of Action Development, Reaction, & Follow-up to Measure Effectiveness

- For DoD this means ***militarily useful reaction*** in ***militarily useful time***
  - Reaction generally must consider mission context, since reaction will be a risk decision

Interoperability in Detection, Diagnosis, & Follow-up

- **Speed is essential, machine-to-machine where we can**
- Detection may be collaborative
- *Diagnosis very often will be*
  - Within DoD & with all those mission partners, industry, etc.
- Follow-up to ensure the course of action selected actually worked may be collaborative
  - Much measuring, sharing, calculating of effect
- **Interoperability for content, identity of devices essential else correlation of data too hard to automate**

Course of Action Development

- What's a COA look like?
  - **Might be a policy change (security, network, etc.)**
  - Local, Global
- **Might have a playbook of closely held policy changes** that anticipate particular mission situations (pre-built COAs)
- When the diagnosis points to a particular situation, execute by deploying new policy (then measuring)
- **All of the policy-based management techniques (application, network, security) will likely be involved in a reaction**

An Example Reaction That Must Be Machine-to-Machine

- "Trust relationships" are great for sharing
- Parties decide to trust each other. To enable better sharing, they become more vulnerable to each other
- This works very well until one of the parties is successfully cyber attacked, and so **in an instant, becomes unworthy of the other partners' trust**
- Detecting that this has happened, and reacting in a mission-useful way is very much an automation challenge

So We've Reduce Attack Surface Automatically, Detected & Reacted to Cyber Attacks Automatically…Are We Done?

Nope.

Achieved a sort of static policy compliance, with reaction to those (few) attacks that breach our protections

Attacks are more expensive, but still cheap relative to other forms of warfare, relatively easy to develop, etc.

We've eliminated whole classes of bad guy, but capable, persistent bad guys will eventually succeed (although we may contain them better)

We may be able to use the automation we've deployed to take the next step

- Drive up cost, time-to-develop, complexity, uncertainty, & risk for adversaries

- Drive down likelihood of (short term or longer term) adversary success

- **Maneuver our infrastructure, maneuver policy faster than an adversary. Be more agile than an adversary**

Security Automation to Change the Infrastructure Faster Than An Adversary Can Develop and Deploy Attacks

Dynamic Defenses

What will this future look like?

- The bad guys are forced to turn to automation which may cause us to change our policy maneuver strategy to one that is more reactive again?

- We end up in a kind of evolutionary algorithm war with bad guys' and defenders' systems evolving ever faster to defeat each other?

- Maybe we get to truly trustworthy systems?

Bad Guys (2)

I Worry A Bit About Over Centralization and Fragility

We Want and Need *Speed* Across Many Devices: Centralized Control of Policy in Many Devices is Like a Phase Change

Something starts it, and it spreads like wildfire, or a shockwave

"I want policy 'RED'"

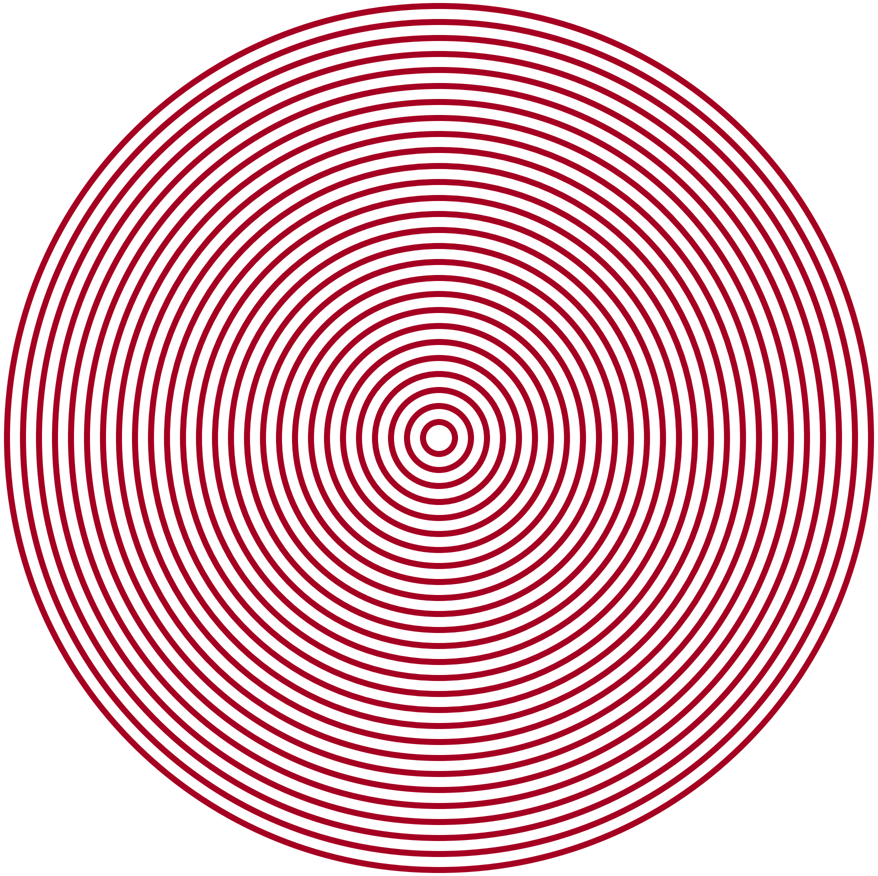

Bad Guy, “But I want policy ‘BLUE’”

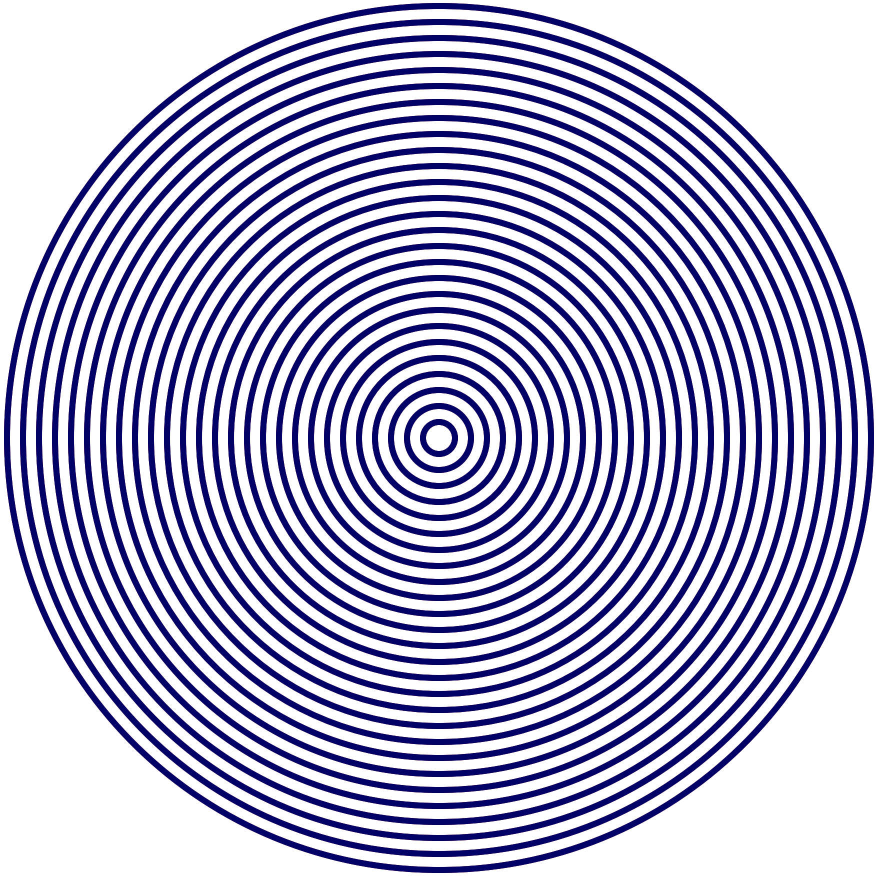

This Whole Ecosystem of Content, Policy Deployment, Measurement Must Be Built *Starting Now* to Have Low Attack Surface

- Strong integrity protection and verification *at exactly the right points* in the production, consumption processes
- Strong interoperable notion of information producer & consumer identity
- Tools that are hard to attack
- Standards for evaluation of components
- Etc.

We Must Also Develop Methods to Detect Bad Guys Infiltrating Our Ecosystem or Subverting Our Tools

(I don't have easy answers for this)

Summary

- **Standards-based, interoperable, unique device identity**
- Ditto for any other other necessary identities
- End-to-end policy construction/analysis, deployment, measurement across complicated infrastructures (including topology/relationships)
- Play-books of militarily useful COAs (again across groups)
- Maneuver faster than the bad guys
- **Do all of this in a realistic operational environment, right alongside those bad guys**

DEFENSE INFORMATION SYSTEMS AGENCY