A Worst-Case Worm

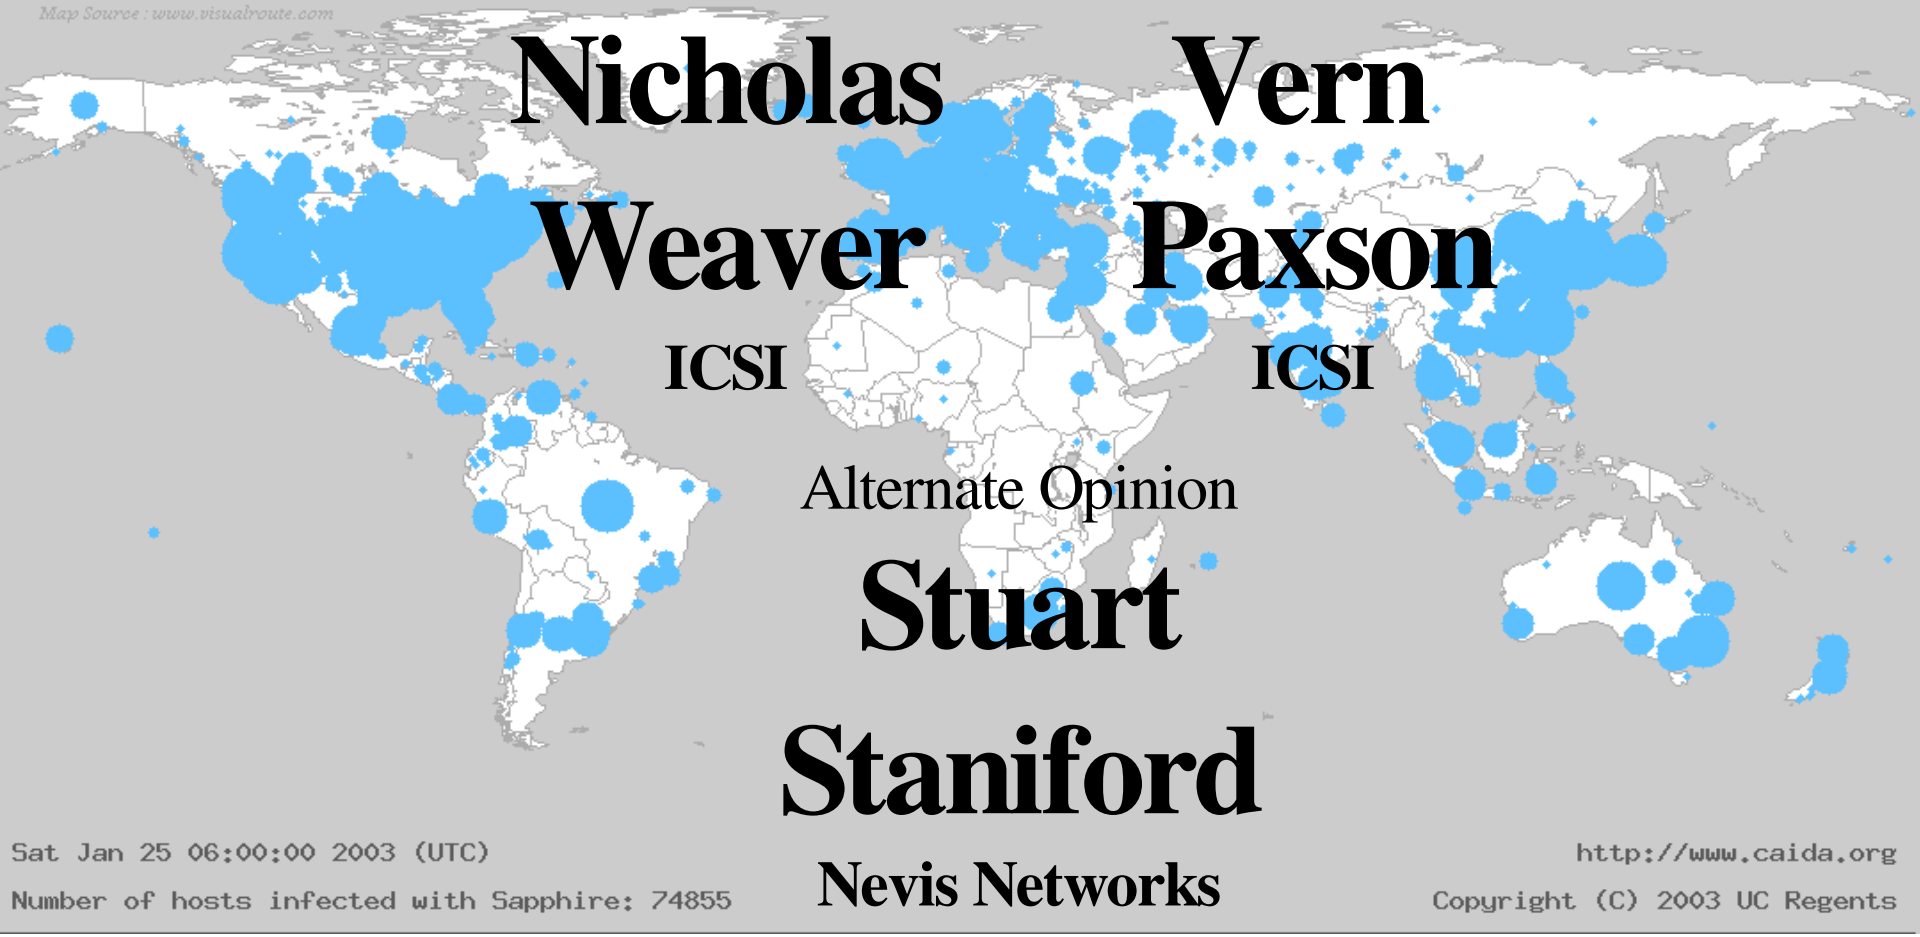

Nicholas Weaver
ICSI

Vern Paxson
ICSI

Alternate Opinion

Stuart Staniford
Nevis Networks

What is our Worst-Case Worm?

- We desire to understand the worst-case possible event
  - We need a defensible estimate
  - Needed to understand the need and requirements for worm-defense: how big is the threat? How much to spend? (if any?)
- What would it look like?
  - Nation-state level adversary
  - "0-day" CIFS/SMB or RPC worm
    - No patches available
  - 0-day firewall crossing routines
  - Malicious payload:
    - Overwrite the disk/remote disks (starting immediately)
    - Reflash the BIOS when possible
    - Infectious reserve to enable reinfection
  - More details in the paper
- We have a good feel for what can be done...

  But are we missing some more imaginative payloads?

But what would this do?

- Need some way of estimating the damage
  - "This would be bad" is insufficient:
    we need to offer reasonable dollar amounts
  - We need transparent estimates, allow others to evaluate our assumptions
    - Please help refine our assumptions and model
  - Understanding the effect of assumptions
- Should we worry?
  - Are the consequences severe enough to justify special efforts at defense?
  - Is this "well founded" FUD?
    - How are our claims different from mi2g saying
      "Bagel, MyDoom, and Netsky together have caused >$100B in damage"
    - Especially since we are not disinterested parties

Overall Damage Model

$$D_{total} = N_{inf} \cdot D_{system}$$

Total Damage — Number Infected * Damage per system

$$N_{inf} = P_{penetration} \cdot N_{vulnerable}$$

Number Infected = Probability of Penetration * Number vulnerable

Recovery Cost + Lost Productivity * Cost of lost productivity

$$D_{system} = D_{rec} + T_{time} \cdot D_{time} + P_{data} \cdot D_{data} + P_{bios} \cdot D_{bios}$$

Probability of Data Loss * Cost + Probability of BIOS flash * Cost

Major Limitations

- Linear and uniform model:
  - Every affected system is of equal importance
    - Averaging only works when most systems are affected
  - Lost System → Disrupted Worker
- Depends greatly on assumed values
  - Are our assumptions reasonable?
- Ignores huge effects:
  - What about worldwide effects?
  - What about critical infrastructure? Would there be any?
  - Do more systems down cause amplification of damage?
  - What about secondary effects?
    - Impossible to estimate
      prone to gross exaggeration
  - What about human ingenuity and adaptation?

Estimating Systems Affected

- ~85 million business PCs
  - Based on a single survey
  - We desire a better number
- Penetration factor:
  - Not all systems will be infected
    - Good firewalls
    - Luck
  - Ignores nonlinear factor:
    - Large institutions are more likely to be affected
  - We assume ~60% penetration
- Thus $N_{affected} = 50$ million

Estimating Recovery Cost

- System administrator time to recover each infected machine
- Large institutions should have mass-install procedures
  - Otherwise, Windows can be difficult to manage
  - Thus system recovery is generally quick per machine
- Smaller institutions will be slower
  - Digging up CDs, swapping disks, more work per machine
    - But still somewhat parallelizeable
- Already understood mechanisms to bring up systems in a hostile environment
  - Blaster/Welchia/Sasser taught us how
- We assume that most will be fairly fast
  - Using ghosting/mass install techniques
  - Average time of ~ ½ hr per system
  - D_{rec} = $20/system

How Much Does a Lost Hour Cost?

- Approximate based on US GDP:
  - US GDP: $11 trillion
    US worker population: 138 million
    Hours Worked per Worker: 34 hours
  - Productivity of a worker-hour: ~$45/hr
- But we should reduce this figure
  - An hour of lost computer time is not an hour of lost productivity
    - Other things can be done when the system is down
    - This is nonlinear: a lost computer-hour is significantly less important than a lost day
  - So we approximate a lost hour (D_{time}) = $35/hr
    - But is this still overestimating the value?

How Much Lost Productivity?

- For a typical system:
  - Data not permanently lost
  - BIOS not permanently corrupted
- We assume 2 work-days of lost time
  - 1 day for Microsoft to develop a patch
    - Not much sleep in Redmond
  - 1 day for the system administrators to recover most systems
    - Not much sleep for the IT staff either
- Thus, our approximation:
  - $D_{time} * T_{time} = \$35/hr * 16hr = \$560/system$

Effect of Lost Data?

- Only one survey we could find
  - D_{data} = $2000 ***if*** the data is unrecoverable
    - and that's a lost data *item*, not a disk
    - Other components from this survey are accounted for in the rest of our model
- What is the probability of unrecoverable data loss?
  - Survey says 20%, but...
  - What is an "incident"?
    - Would 4 systems in the same institution be an incident?
  - We err on the conservative side and assume 10%
  - Where do they get $2000 anyway?
- Thus, our approximation:
  - $D_{data} * P_{data}$ = $2000 * 0.1 = $200/system

Effect of Reflashed BIOSes?

- The BIOS (non-volatile-memory) is commonly soldered to the motherboard
  - Used to provide the initial program to boot the system
  - Many systems have no recovery procedure if the BIOS is corrupted
    - Others have a jumper to allow recovery
    - Others are socketed for an easy exchange
  - Corrupting the BIOS is usually a vendor/motherboard specific action
    - The video-card BIOS may also be vulnerable
- This is a contentious area:
  - Vern and I believe the damage is only moderate from this attack
  - Stuart thinks we are substantially underestimating the impact

Our Assumed Damage From BIOS Reflashing

- Probability of a BIOS attack succeeding?
  - Vendor specific routines required
  - Not all BIOSes can be permanently corrupted
    - Attacker will select those which maximize the damage
  - We assume $P_{bios} = 0.1$
- What is the damage?
  - New system/external costs -> $1000/system
    - Grabbing new machines/paying for recovery services
  - Additional lost productivity of a week -> $1400
- Thus our approximation:
  - $P_{bios} * D_{bios} = .1 * \$2400 = \$240$

Overall Damage and Effect of Assumptions

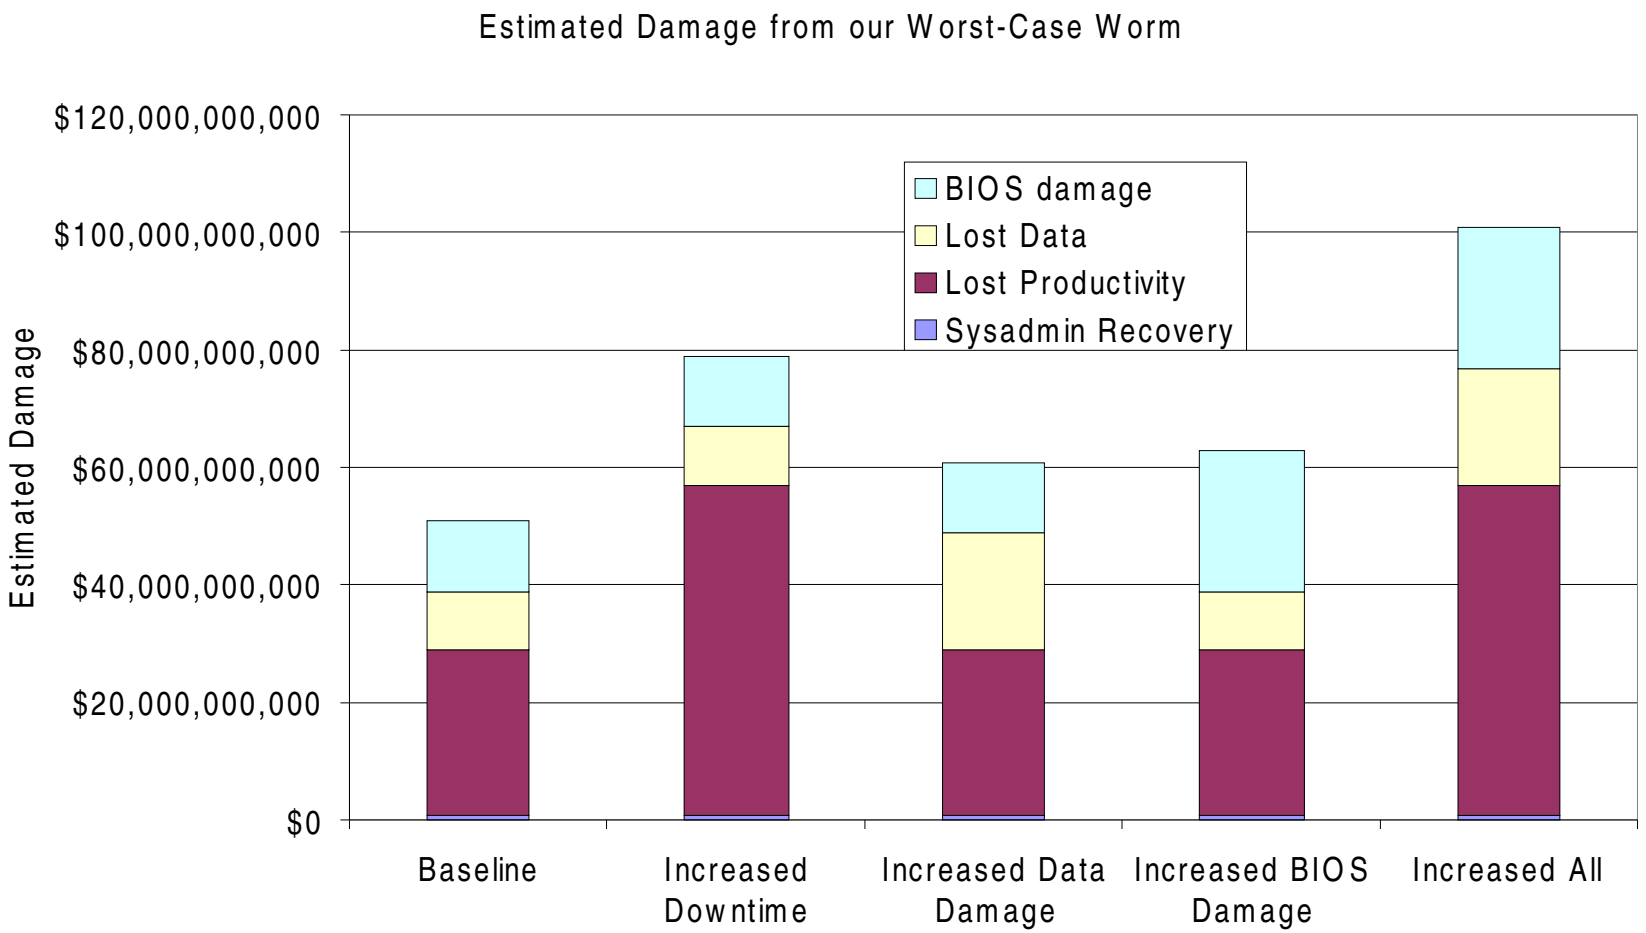

Primary Conclusions

- This appears significant
  - $50 billion dollars is a large amount of damage
  - This is worth worrying about
    - We should probably build defenses now:
      Microsoft Windows ***is*** Critical Infrastructure
- Excluded factors and assumptions are problematic...
  - What would happen to critical infrastructure?
  - Are computers really ***this*** important?
- "Weapons of Mass Annoyance" are a concern
  - Enough annoyance becomes significant

An Area of Debate: BIOS Damage?

- Could it be 30% affected? 90%+?
  Would downtime be a week? Or More?
- Nick & Vern's View:
  - Our figures are reasonably realistic
    - Significant downtime probably can't extend much beyond a week
  - But when in doubt, bias towards conservative
- Stuart's View:
  - We are substantially underestimating the impact of a BIOS reflashing attack

Nick & Vern's Arguments:

- Our model's assumptions are already being pushed by assuming >1 week's downtime in some cases
  - More important systems will be recovered faster through market forces, violating a basic assumption
- Humans are incredibly adaptable
  - Product cycles and procedures go out the window
    - All Nighters and Mountain Dew
  - There appear to be recovery kludges
    - Significant recovery could occur in 2-4 days even in the worst case
- As a result, we might even be ***overestimating*** the impact of a BIOS reflashing attack
  - But in our current model, the significance is secondary

Important Systems and Market Forces

- Some systems are nearly valueless:
  - Old "boat anchors" in various offices
- Some systems are highly valuable
- Once there are a significant number of operational systems, market forces (internal and external) will reallocate them
  - Take systems from home
  - Buy used but working systems
- Thus although many systems might be down for weeks
  - These will be nearly valueless systems
  - Our model can't cope with this, as we assume all systems are the same
- Assuming 1 week may be too long, as nonfuctional valuable systems are swapped with the working valueless
  - As soon as >25% are recovered from BIOS damage, we can probably assume a large fraction of the disruption has passed

Human Adaptation: Repairing Motherboards?

- Rework stations
  - Desolder and resolder a replacement
    - ~15 minute operation for a skilled user
    - Could easily charge $200-1000 per motherboard
- In-situ BIOS reflasher
  - Probe card attaches to the BIOS pins
  - Attaches to a notebook to download the new BIOS image
  - Outcompetes the chipset to provide the new data
  - Looks like a 24-hour hack to come up with a crude design
    - Crisis situation only needs 80% success
    - Items purchasable from Fry's
  - Once designed, can be distributed & fabricated
    - Hand information to Dell etc...
- Thus recovery looks faster than 1 week... Oops...

However...

- 50M systems may be an underestimate/overestimate...
  - How many ***business significant*** computers are there?
    - We only really want to count significant systems, not all systems
    - Who can measure this?
- Are we missing some other damage modes?
  - Would capturing access to 10,000 (100,000?) brokerage acounts with $50,000 each allow more damage?
    - It would certainly allow easier/more profit...
- Does it matter?
  - If the answer is $50B, $100B, or $500B, does this change our view of how much to spend on defenses?
    - What if it's just $10B?
  - If we ***say*** $50B, $100B, or $500B, does this change how others perceive our work?
- The worm crowd says "Fix the Economics"
  - The propigation/attack seems very plausible