# Trustworthy Software

Jeannette M. Wing

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA  USA

Grand Challenges in Trustworthy Computing
CRA Conference at Snowbird
13 July 2004

# Disclosures and Apologies

- One-year sabbatical at Microsoft Research, working on security. Organized MSR-UW-CMU workshop on Software Security.
- Member of Microsoft Trustworthy Computing Academic Advisory Board
- Member of NAS Computer Science and Telecommunications Board study IT subpanel for *Making the Nation Safer: The Role of Science and Technology in Countering Terrorism*.

- CRA Rule: No individual accreditation allowed.
  – I will be acknowledging individuals, but they were not at the CRA meeting.

# Starting Point: Secure Software

- An Ongoing Challenge: Security
  - It’s about software, not the network.

- An Ongoing Challenge: Software Engineering
  - Forget trying to solve the general problem.
  - Solve it for one class of properties.
  - Choose that class today to be one that is critical, timely, and of societal benefit.
  - For example, security!

# A Grandiose Goal: Trustworthy Software

- Trustworthy =
  - Reliability
    - Does it do the right thing?
  - Security
    - How vulnerable is it to attack?
  - Privacy
    - Does it protect a person's identity?
  - Usability
    - Can a human use it easily?

- Focus here on software, but could broaden to "computing" or "systems."

# What is Reliability?  Security wrt Reliability?

- Reliability
  - Formally, the system meets its spec (focus on correctness)
    - In theory, we know how to do this
    - In practice, it doesn't scale
    - And worse, we can't get the formal spec right
      - Though we can get partial specs right, e.g., types
  - Informally, users aren't surprised
    - Hence, depends on user expectations
    - "Good enough" for need, e.g., land-line vs. cell phones (Lampson)
- Security
  - Goal: Stop people from doing bad things
  - In principle, security spec can be part of reliability spec (see above)
  - In practice
    - Fault/threat models are dramatically different
      - Cannot estimate the probability of attack
    - Designing for reliability is different from that for security
      - Redundancy vs. diversity
      - Separation of abnormal modes from normal modes vs. failure of "Trusted Computing Base"

# Reliability and Security

Idea #1: Correctness, but with an attacker in mind

- Pre-conditions, or more generally, assumptions about the software's environment, tell the attacker exactly how to break your system.
- E.g., a well-typed program is type-correct wrt the (implicit) assumption that the type checker is correct.

Idea #2: "Good enough" vs. absolute

- "Good enough" is determined by user, e.g., land lines vs. cell phones (Lampson)
- "Good enough" is as much as we can obtain in practice.

Idea #3: Quantifiable metrics for security as we have for reliability, e.g.,

- Economics/financial model (CRA Grand Challenge #4)
- Underwriters lab model
- Game theoretic model (payoff and reward functions)

Idea #4: Looking above the level of code, beyond buffer overruns

- Anticipate tomorrow's attacks
- Look for vulnerabilities at design and architectural levels of software

# Secure By Design: What We Need

- Compositional techniques
  - To discover interface mismatches, e.g., DNS+Netscape vulnerability
  - To detect clashing security policies, e.g., IE and Outlook settings
  - To anticipate emergent abusive behavior, e.g., spam, Google bombs
    - Udi Manber's penny-change box analogy

- Design principles
  - Security design principles with software in mind
    - E.g., Defense in Depth, Principle of Least Privilege, Secure by Default
  - Software design principles with security in mind
    - E.g., Weaken/check pre-condition, strengthen post-condition, document invariants and abstraction functions
  - Something akin to Abadi and Needham's crypto protocol design principles

# Secure by Design: MS03-007 Windows Server 2003 Unaffected example from David Aucsmith

| | | Defense in Depth |
|---|---|---|
| The underlying DLL (NTDLL.DLL) was not vulnerable | Code made more conservative during the Security Push | Check Precondition |
| *Even* if it was vulnerable | IIS 6.0 not running by default on Windows Server 2003 | Secure by Default |
| *Even* if it was running | IIS 6.0 doesn’t have WebDAV enabled by default | Secure by Default |
| *Even* if it did have WebDAV enabled | Maximum URL length in IIS 6.0 is 16KB by default (> 64KB needed for exploit) | Tighten precondition, Secure by Default |
| *Even* if the buffer was large enough | Process halts rather than executes malicious code, due to buffer overrun detection code (-GS) | Tighten Postcondition, Check Precondition |
| *Even* if there was an exploitable buffer overrun | Would have occurred in *w3wp.exe* which is now running as `network service’ | Least Privilege |

# Privacy: Questions to Ponder

1. What does privacy mean?

2. How do you state a privacy policy? How can you prove your software satisfies it?

3. How do you reason about privacy? How do you resolve conflicts among different privacy policies?

4. Are there things that are impossible to achieve wrt some definition of privacy?

5. How do you implement practical mechanisms to enforce different privacy policies? As they change over time?

6. How do you measure privacy?

# What is Privacy? Societal Answers –thanks to Doug Tygar

- Fundamental human right (European Council)
  - Problem: People don’t treat it like a fundamental human right.  E.g., people don’t sell right of freedom of religion for 20 cents off olives.
- Property right (many, many libertarian geeks)
  - Problem: No negotiation power, no way to take back property
- “Penumbra of the constitution” (Supreme Court, Roe v. Wade)
  - Problem: Widely criticized
- Subject to regulatory constraints (current US practice)
  - Problem: Patchwork approach to privacy, i.e., Bork Bill (for video records)
- “You have no privacy, get over it.” (Scott McNealy)
  - Problem: People do seem to desire privacy
- Spy states are good for you (David Brin)
  - Problem: But everybody should be part of the action

# What is Privacy? Technical Answers –thanks to Doug Tygar

- Privacy is like confidentiality
  - Need to specify who has access to private data
  - Want to prevent data from accidental disclosure
- Privacy is like DRM
  - We release information to an outside party, want to restrict its use
  - Should watermark data
  - Possibility for NGSCB (formerly Palladium) or hardware-based protection
- Privacy is like data mining
  - Information can be released in statistical summary or in individual parts
  - Question of whether sensitive information can be inferred
- Privacy is like data escrow
  - Information is recorded, and the fact that it is recorded is known
  - Need a “search warrant” or other authorization to access information

# Usability

- Seminal paper: Why Johnny Can’t Encrypt? – Whitten and Tygar, Usenix’99
- Balance between usability and security, privacy, and reliability
  - Password standards vs. writing your password down
  - XP “Phone Home” automatic updates and bug fixes
  - Microsoft Watson bug database
- How much does the user need to know? How much control should the user have?

# Clicking Your Way Through Security

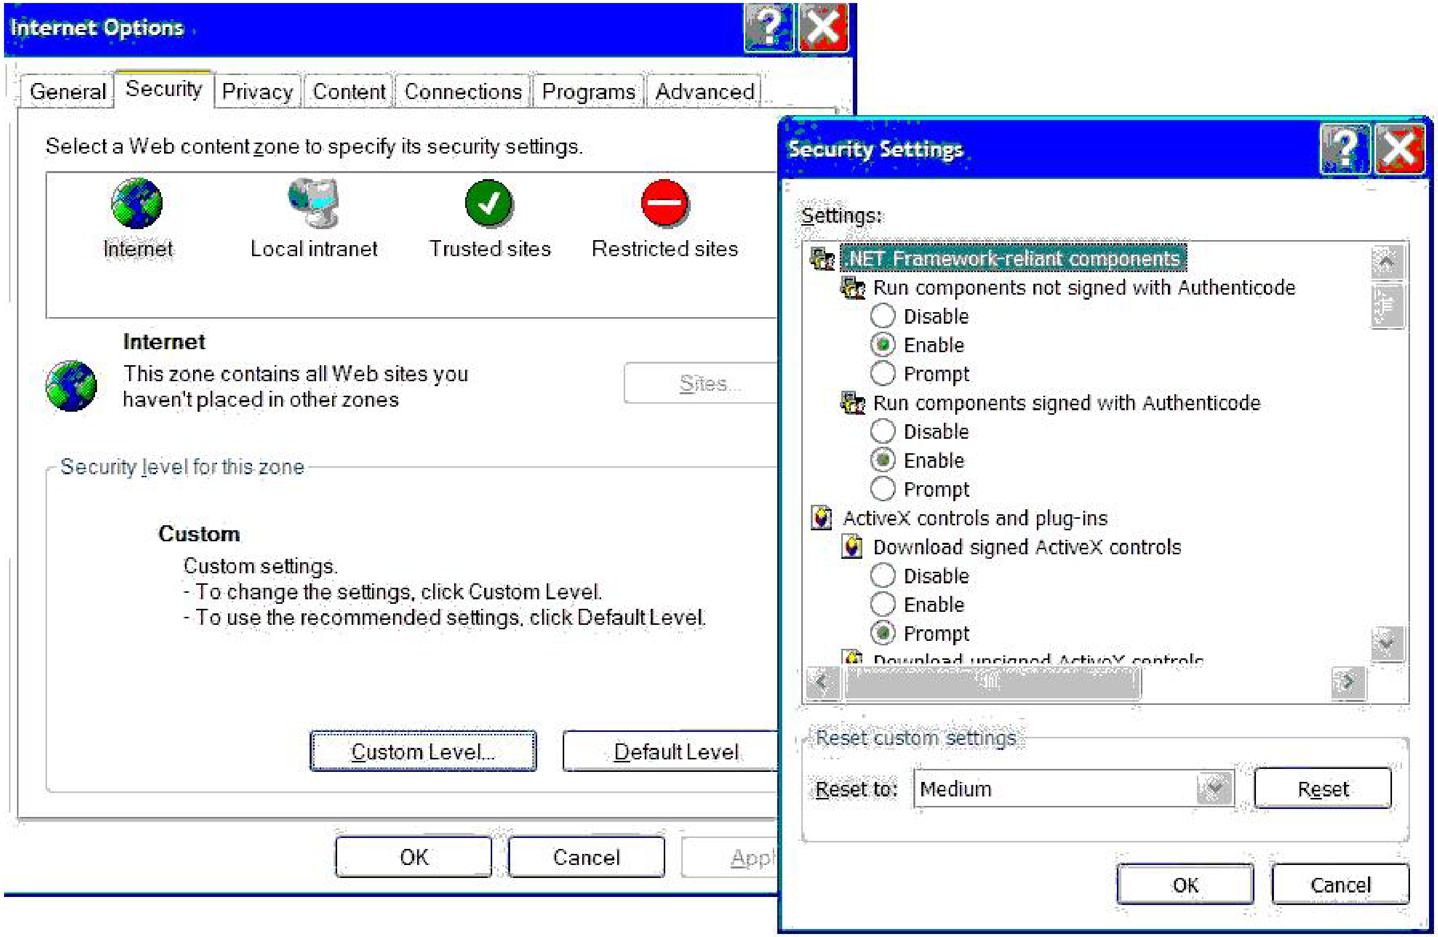

Select a Web content zone to specify its security settings.

**Internet**

This zone contains all Web sites you haven't placed in other zones

Security level for this zone

**Custom**

Custom settings.
- To change the settings, click Custom Level.
- To use the recommended settings, click Default Level.

Settings:

.NET Framework-reliant components
- Run components not signed with Authenticode
  - ○ Disable
  - ● Enable
  - ○ Prompt
- Run components signed with Authenticode
  - ○ Disable
  - ● Enable
  - ○ Prompt

ActiveX controls and plug-ins
- Download signed ActiveX controls
  - ○ Disable
  - ○ Enable
  - ● Prompt
- Download unsigned ActiveX controls

Reset custom settings

Reset to: Medium

# Do You Read These?  What Are They Saying?

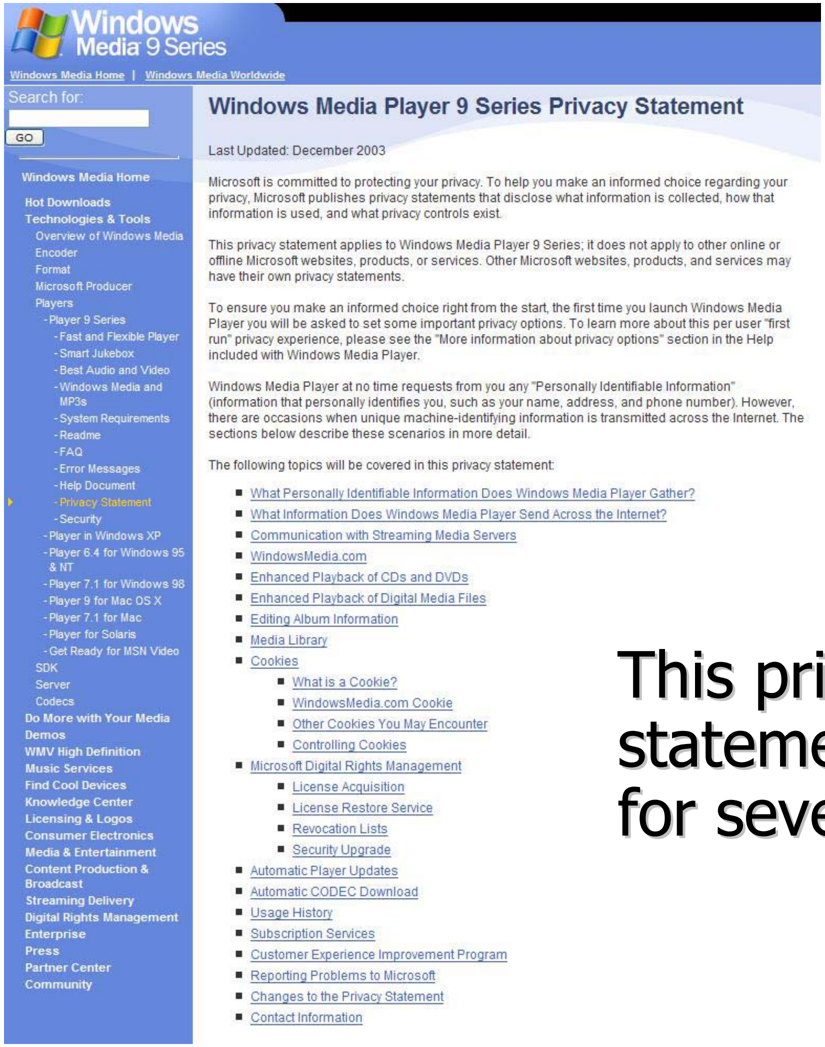

## Windows Media Player 9 Series Privacy Statement

Last Updated: December 2003

Microsoft is committed to protecting your privacy. To help you make an informed choice regarding your privacy, Microsoft publishes privacy statements that disclose what information is collected, how that information is used, and what privacy controls exist.

This privacy statement applies to Windows Media Player 9 Series; it does not apply to other online or offline Microsoft websites, products, or services. Other Microsoft websites, products, and services may have their own privacy statements.

To ensure you make an informed choice right from the start, the first time you launch Windows Media Player you will be asked to set some important privacy options. To learn more about this per user "first run" privacy experience, please see the "More information about privacy options" section in the Help included with Windows Media Player.

Windows Media Player at no time requests from you any "Personally Identifiable Information" (information that personally identifies you, such as your name, address, and phone number). However, there are occasions when unique machine-identifying information is transmitted across the Internet. The sections below describe these scenarios in more detail.

The following topics will be covered in this privacy statement:

- What Personally Identifiable Information Does Windows Media Player Gather?
- What Information Does Windows Media Player Send Across the Internet?
- Communication with Streaming Media Servers
- WindowsMedia.com
- Enhanced Playback of CDs and DVDs
- Enhanced Playback of Digital Media Files
- Editing Album Information
- Media Library
- Cookies
  - What is a Cookie?
  - WindowsMedia.com Cookie
  - Other Cookies You May Encounter
  - Controlling Cookies
- Microsoft Digital Rights Management
  - License Acquisition
  - License Restore Service
  - Revocation Lists
  - Security Upgrade
- Automatic Player Updates
- Automatic CODEC Download
- Usage History
- Subscription Services
- Customer Experience Improvement Program
- Reporting Problems to Microsoft
- Changes to the Privacy Statement
- Contact Information

This privacy statement goes on for seven screenfuls!

# Trustworthy Software

- Reliability
  - Focus on correctness
  - Goal:
    - Identify ways to check/detect interface mismatches for design-level vulnerabilities.
    - Define compositional ways of reasoning for scalable verification
- Security
  - Focus on authorized access
  - Goal:
    - Identify software design rules/principles with security in mind.
- Privacy
  - Focus on authorized use, perhaps after release
  - Goals:
    - Identify a mathematical structure for privacy analogous to what Lampson's access matrix is for security.
    - Define a logic for reasoning about privacy.
- Usability
  - Humans are often the weakest link.
  - Goal: Balance between convenience and control.

# Why This is Important for Society

- Timely
- What IT (and even non-IT) companies want
- What policymakers and lawyers need
- It’s an international, not national issue
  - E.g., Germany’s privacy laws, globalization of corporations
- Our role as scientists in society

# Trustworthy ~~Security~~ Axiom

Good guys and bad guys are in a never-ending race!

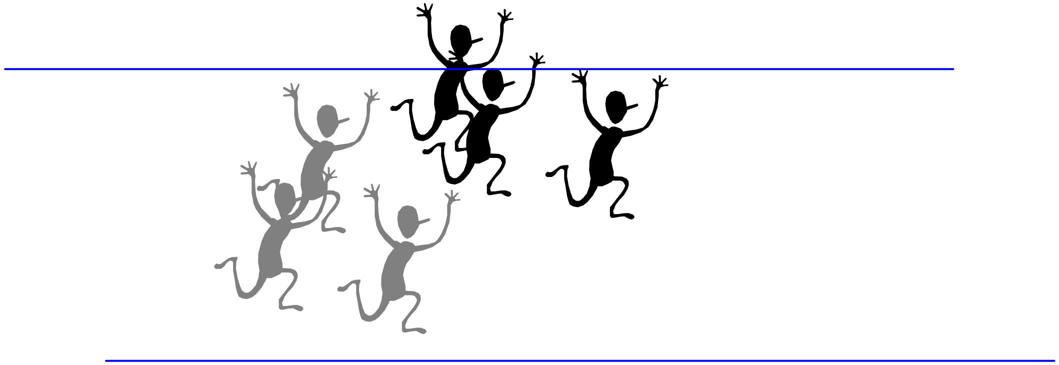

Thus, Trustworthy Software is not really a Grand Challenge, but a Grandiose Goal.