

Resilience and Security

Matt Bishop

Dept. of Computer Science


University of California at Davis

Theme of the Talk

Resilience is holistic not just in terms of the system, but also in terms of the goals of the system.

And without considering those goals, the term “resilience” is **meaningless**

Resilience: Informal Definition



Definitions

- *Resilient*: “resuming the original shape or position after being bent, compressed, or stretched”
- Implication: system may degrade, but will recover
 - Recovery may take quite a while ...

Definitions

- *Robust*: “strong and hardy in body or constitution; possessed of rude strength”; “strong, vigorous, healthy”
- Implication: difficult or impossible to degrade system under attack or through accident
 - May be able to cause loss of functionality, and if so does not imply an ability to restore it

Definitions

- *Survivable*: “capable of continuing to exist after some person, thing, or event; to last on”
- Implication: system will continue to exist, and will continue to function at or above minimum requirements
 - But it may never fully recover

Nuances

- Survivable systems continue to function but possibly at a degraded level
- Robust systems hard to enter compromised state; they may fail instead
- Resilient systems are systems that return to their desired level of functionality

Example: Network Router

Suppose it has a guaranteed minimum QoS.

After attack:

- Survivable with respect to that attack if it will continue to function at least at that level
- Robust with respect to that attack if it will continue to function at the desired level (which may be above minimum one) or fails completely
- Resilient with respect to that attack if it recovers and functions at desired levels; it does not fail permanently

Two Common Aspects

- Time needed for the system to return from a perturbed state to an equilibrium (Pimm, 1984)
- Maximum perturbation that will *not* prevent the system from returning to equilibrium (Holling, 1973)

Dimensions of Resilience

- Availability: the most studied
- Integrity: much harder
- Confidentiality: seems to be the hardest

Availability Resilience

- Discussed widely in the literature
- Many metrics:
 - Recovery time
 - Number of functioning components vs. total number of components
 - Effect on network bandwidth vs. capacity
 - Ability to adapt existing capabilities to new requirements
- Metrics must be chosen so that they apply to the mission of the system

Integrity Resilience

- Usual approach is to replicate data
 - If replicas disagree, apply a Byzantine or voting algorithm to determine correct values
 - But now data is more available to adversaries
 - Also, probably a monoculture ...
- Suppose data is compromised
 - Generally care about outputs
 - Often outputs can be “close enough”; goal is to avoid those that cause damage

Key Ideas

Holistic approaches

- Systems must be designed so that a single compromised component cannot produce a damaging output
 - Physical, virtual separation of components
- Outputs must “make sense” in the context in which they are to be used
 - Think “back of the envelope” calculations

State of the Art

Focus is on first and not second

- *N*-version programming
- Swarm systems
 - Resilience tied to the swarm, not to individual elements of it

Origin Integrity Resilience

- What happens if a source is mis-identified or the original input is bad?
 - If values or trustworthiness of outputs interpreted in light of origin ...
 - Example: provenance; what happens if the data put into the provenance is not accurate?

Assurance and Resilience

- Trust is bound to assurance
 - “Assurance” is confidence that an entity meets its requirements, based on specific evidence obtained by applying specific techniques
- What happens when assurance is compromised
 - How do you regain the desired level of assurance?
 - How do you tell when you have regained it?

Integrity Challenges

- How can we evaluate goodness (adequacy) of outputs in the context of use?
 - How to determine the relevant aspects of context
 - How to determine a methodology for doing this
- Once assurance is compromised, how does one regain the desire level of assurance — and how does one *know* that level has been regained?

Confidentiality Resilience

- Harder than the other two types
- Goal is to *regain* confidentiality in some form
 - Delete information from the Internet?
 - Hide information
 - Change the context so it applies to something or someone else
 - Bind it to something the recipient of the leak wants to stay hidden

Beginnings

- Shannon: information transmission is about uncertainty
 - Level 1: How accurately can the symbols of communication be transmitted? (technical)
 - Level 2: How precisely do the transmitted symbols convey the desired meaning? (semantic)
 - Level 3: How effectively does the received meaning affect the conduct in the desired way? (effectiveness)

Hiding Information

- Flooding: bury the data to be concealed in lots of non-confidential (or seemingly confidential) data
- Deception: give the adversary a combination of correct and incorrect but believable data

Context is Critical!

- Flooding assumes I can't distinguish relevant from irrelevant information
- Disseminating false information assumes it can't all be checked
 - ... at least not in the time period of interest

Uncertainty's Effectiveness

Uncertainty useful only if it inhibits undesirable actions

- Suppose I know that one person in a group of k has an uninsurable disease
 - If I don't know which one, I deny them all insurance

False Attribution

- Am I Matt Bishop the UCD professor or the UCL professor?
 - ... or on the faculty at both places?
- Solution: identity management
 - How do you verify identity initially?
 - How do you handle those not in the system?
 - How do you ensure the system is up to date?

Changing Semantics

- Teacher denied promotion because a picture of her as a “drunken pirate” found on Internet
 - No proof she was drinking anything alcoholic
 - Change caption to “Trying homemade concoction of various juices”
- London court: “How could someone kill a baby?”
 - Stress, intonation convey interpretation

Defeat Use

- Newspaper asked for list of political donors
 - Candidate refused to provide it to prevent donors from being approached by others
 - Legal, back then
- Newspaper ran story suggesting candidate had something to hide
- Candidate sent list to reporter, but said he could print it only if he printed **all** the names
 - Publisher of newspaper had donated a lot

Composition Problem

- Information set X has no confidential information
- Information set Y has no confidential information
- Information set $X \cup Y$ has very confidential information
- Because of a change over which you have control, information set X now contains information that should have been confidential

Confidentiality Challenge

- How can we evaluate what data is to be kept confidential?
 - Problems arise with temporal, composition issues
- Once confidential data is leaked, can the system be reconfigured or updated to function in such a way that the leaked information need no longer be kept confidential?

Metric Challenge

- Develop non-binary approaches to measuring data confidentiality, integrity
 - Some work on probabilistic approaches to measuring data integrity exist, but is scant
- Then extend them to measuring resilience of confidentiality, integrity
 - We'll get to metrics later ...

Resilience of What?

Approaches differ depending on what you want to be resilient, which is driven by goals or requirements

- Information
- Entities
 - Systems
 - Networks
 - Other Resources

Resilience of Information

- The information must be sufficiently correct and trustworthy to be used as needed (desired)
- Independent of systems or other entities

Resilience of Entities

- The entity must perform its desired function in the face of obstacles (attacks, failures, etc.)
- Usually seen as availability
- Integrity here very important
 - Relates to trustworthiness
- Confidentiality depends on context
 - You may want to hide the *presence* of the entity
 - You may want to hide the *absence* of the entity

Metrics for Resilience

- Assume:
 - System provides well defined set of services
 - Robustness, survivability, resilience defined for it
 - Services have requirements that define quality
- Thought: quality of service may measure aspects of confidentiality, integrity as well as availability
 - Quality (level) of assurance
 - Quality of encryption (key length, cryptosystem)

R₄ Framework

- *Robustness*—the ability of systems, system elements, and other units of analysis to withstand disaster forces without significant degradation or loss of performance;
- *Redundancy*—the extent to which systems, system elements, or other units are substitutable, that is, capable of satisfying functional requirements, if significant degradation or loss of functionality occurs;
- *Resourcefulness*—the ability to diagnose and prioritize problems and to initiate solutions by identifying and mobilizing material, monetary, informational, techno-logical, and human resources; and
- *Rapidity*—the capacity to restore functionality in a timely way, containing losses and avoiding disruptions.

System Aspects

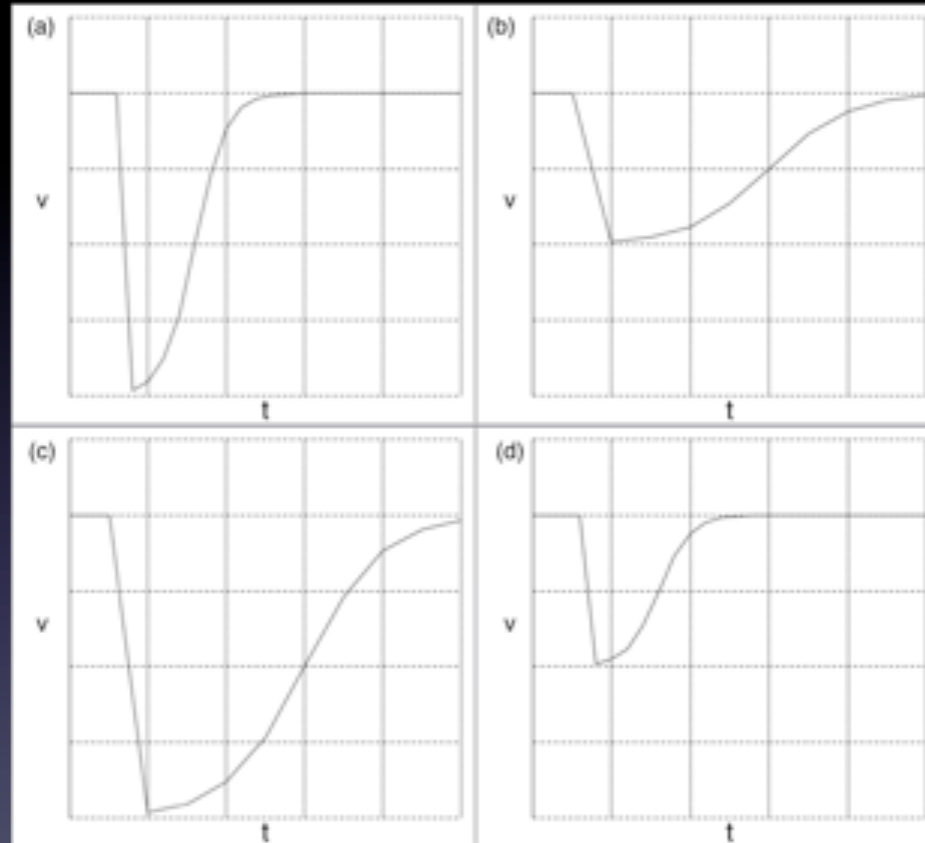
Resilience has 3 aspects

- Reduced probability of failures
- Reduced consequences from failures
- Reduced time to recovery

Vector Measurements

- Measurements are vectors (may be collapsed to a number)
 - Many different components to resilience (see R₄ for an example), so need to measure each one
 - When collapsed, need to take into account how each component of resilience affects the others

Example



From Ford, Carvalho, Mayron, Bishop, "Towards Metrics for Cyber Resilience,"
21st EICAR Annual Conference Proceedings pp. 151–159 (May 2012).

Perturbations

- Measurements are affected differently by different perturbations
 - So the failure and subsequent recovery, or the compensation, will vary and probably cause differences in outputs
 - Measurements measure resilience of system with respect to one or more particular perturbations

Don't Assume Randomness

- Metrics must handle both random and deliberate, sustained problems that impact functionality
 - For random errors, an average-case model of measurement works
 - For sustained errors, a worst-case model of measurement is needed

Example

- 10 servers fail due to random events, 9 of them fixed
 - Risk greatly diminished as 9/10 are working
- Same 10 servers fail because an adversary exploited a vulnerability, 9 of them patched and restored
 - Risk still there as attacker need only locate that 1 unpatched server

What Is Your System?

- Metrics depend on what the system is defined to be
 - If a component is resilient and a perturbation occurs, component will recover
 - If other components depend on that component, metric must include them in resilience measurement

Delta Air Lines Failure

- 2016 malfunction in “a power control module [caused] a surge to the transformer and a loss of power”
 - Power quickly restored — metric would show that aspect was (somewhat) resilient
- “Some critical systems and equipment didn’t switch over to backups, causing ‘instability’”
 - Metric for system as a whole would show much lower resilience

Implication #1

- Ranking of systems being compared for resilience varies depending on context (dependencies, environment) and use (application)
 - Delta requires its systems to serve customers, so one metric is time until customer service restored to previous level
 - Power company wants to restore service to the data center, so its metric is time until the generators produce power at previous levels

Implication #2

- Combining multiple resilience metrics and scenarios to produce a global total ordering of entities must be done on a case-by-case basis
 - This is due to context
 - Problem is using a scalar rather than a vector of metrics

It's Not Just Outputs

- Metrics must take into account more than outputs
 - Web server example
 - Capacity to handle connections has dropped (90% of what it was)
 - Moral: take effects of system disruption into account as well as effect on outputs

To Sum Up

- Resilience is holistic with respect to the system
- Resilience is holistic with respect to the overall goals of the system
 - Confidentiality, integrity, availability, performance, usability, ...
- Resilience with respect to confidentiality, integrity needs far more study

Closing Thought: Clear Overall Goals

Gentlemen,

Whilst marching from Portugal to a position which commands the approach to Madrid and the French forces, my officers have been diligently complying with your requests which have been sent by H.M. ship from London to Lisbon and thence by dispatch to our headquarters.

We have enumerated our saddles, bridles, tents and tent poles, and all manner of sundry items for which His Majesty's Government holds me accountable. I have dispatched reports on the character, wit, and spleen of every officer. Each item and every farthing has been accounted for, with two regrettable exceptions for which I beg your indulgence.

Unfortunately the sum of one shilling and ninepence remains unaccounted for in one infantry battalion's petty cash and there has been a hideous confusion as to the number of jars of raspberry jam issued to one cavalry regiment during a sandstorm in western Spain. This reprehensible carelessness may be related to the pressure of circumstance, since we are war with France, a fact which may come as a bit of a surprise to you gentlemen in Whitehall.

This brings me to my present purpose, which is to request elucidation of my instructions from His Majesty's Government so that I may better understand why I am dragging an army over these barren plains. I construe that perforce it must be one of two alternative duties, as given below. I shall pursue either one with the best of my ability, but I cannot do both:

1. To train an army of uniformed British clerks in Spain for the benefit of the accountants and copy-boys in London or perchance:
2. To see to it that the forces of Napoleon are driven out of Spain.

—Duke of Wellington, to the British Foreign Office, London, 1812

Thanks To ...

- Many colleagues
 - Marco Carvalho, Richard Ford, Liam Mayron (resilience)
 - Emily Rine Butler, Kevin Butler, Carrie Gates, Steven Greenspan (confidentiality issues)
- Awards CCF-0905503, CNS-1049738, CNS-1219993, and CNS-1258577 from the National Science Foundation to the University of California at Davis

Any Questions?

- I would prefer answers, but ...