Research Directions in Intrusion Tolerant Systems

Dr. Carl E. Landwehr,
Program Director
National Science Foundation
clandweh@nsf.gov
+1 703-292-8910

IFIP WG 10.4 42nd Meeting

June 29, 2002

2002 ITS Workshop

DSN 2002, June 24, 2002

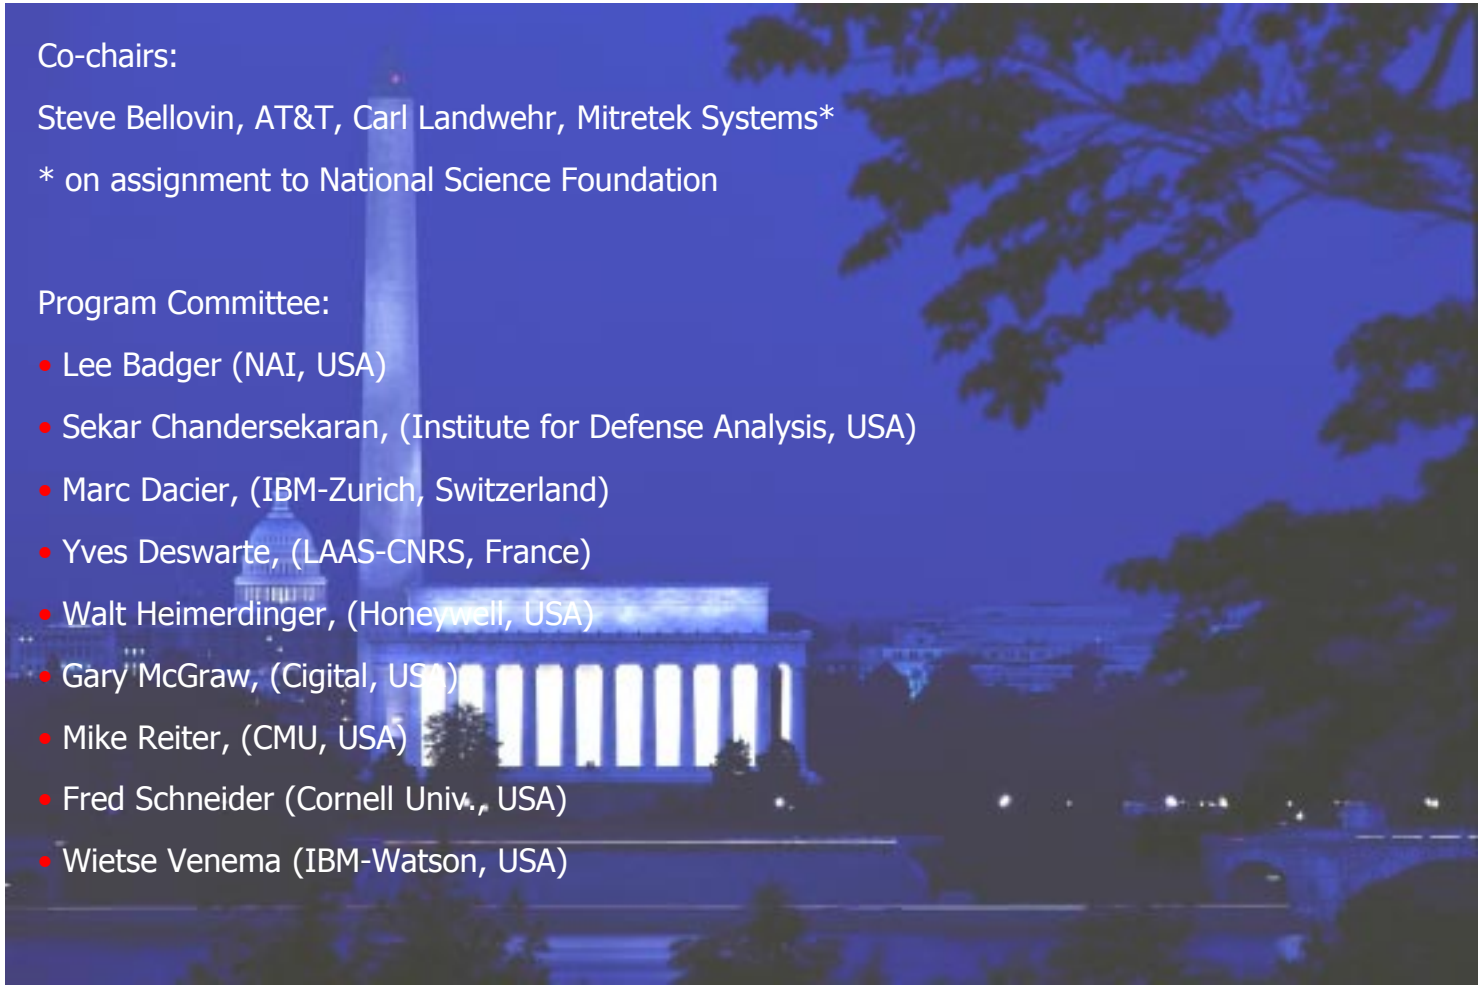

Co-chairs:

Steve Bellovin, AT&T, Carl Landwehr, Mitretek Systems*

* on assignment to National Science Foundation

Program Committee:

- Lee Badger (NAI, USA)
- Sekar Chandersekaran, (Institute for Defense Analysis, USA)
- Marc Dacier, (IBM-Zurich, Switzerland)
- Yves Deswarte, (LAAS-CNRS, France)
- Walt Heimerdinger, (Honeywell, USA)
- Gary McGraw, (Cigital, USA)
- Mike Reiter, (CMU, USA)
- Fred Schneider (Cornell Univ., USA)
- Wietse Venema (IBM-Watson, USA)

Blue Team: ITS Architectures

- **Al Valdes**, SRI International: An Adaptive Intrusion-Tolerant Server Architecture
- **John Knight**, U. of Virginia: The Willow Architecture
- **Franklin Webber**, Contractor to BBN: Intrusion Tolerance through Unpredictable Adaptation (**ITUA**)
- **James Just**, Teknowledge: Hierarchical Adaptive Control for QoS Intrusion Tolerance (**HACQIT**)
- **Peng Liu**, U. of Maryland - Baltimore County: Intrusion Tolerant Database System (**ITDB**)
- **Paul Ezhilchelvan**, U. of Newcastle: A Middleware Architecture for Intrusion- and Fault-Tolerant Service Replication

Each presenter given 10 uninterrupted minutes to introduce his/her system/subsystem. At the end of each presentation, the red team panel given two minutes for specific questions.

Each presentation is to cover, briefly but clearly,
1. Technology description: scope, components, interconnections, operational characteristics
2. Assumptions: what parts of the problem is the project depending on others to solve?
3. Planned reactions to attacks: How are the tolerance mechanisms expected to behave in normal operation and in the face of attacks?
4. What level of degradation is expected in response to particular attacks?
5. Results, if any, of experiments, prototypes

Hierarchical Adaptive Control of QoS for Intrusion Tolerance - HACQIT

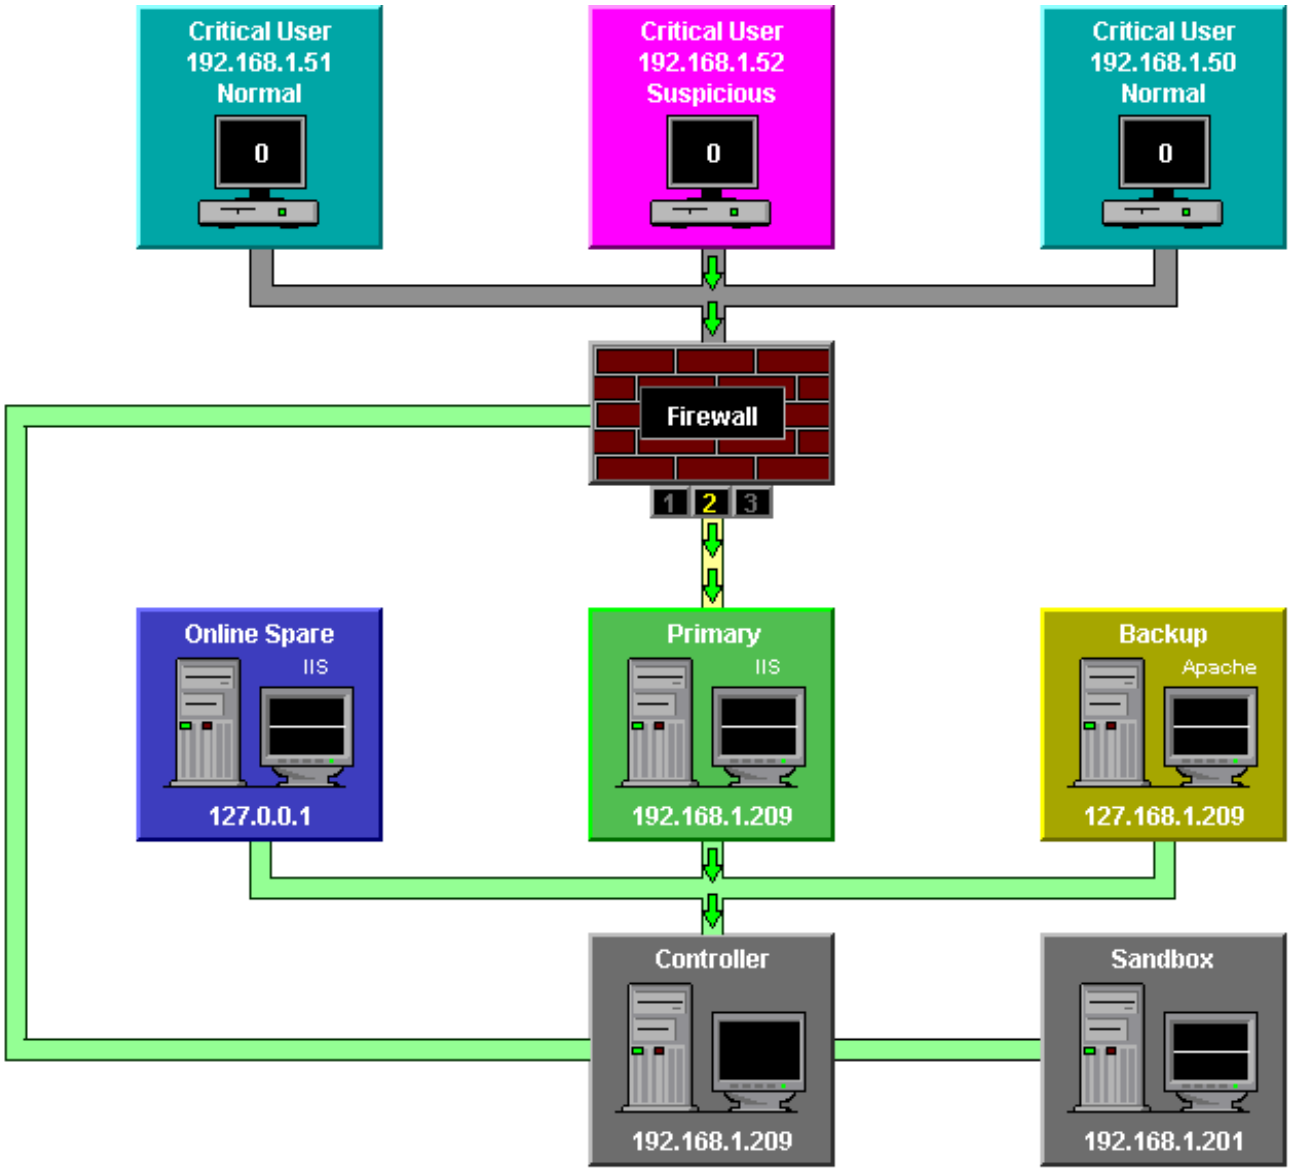

Willow System Architecture

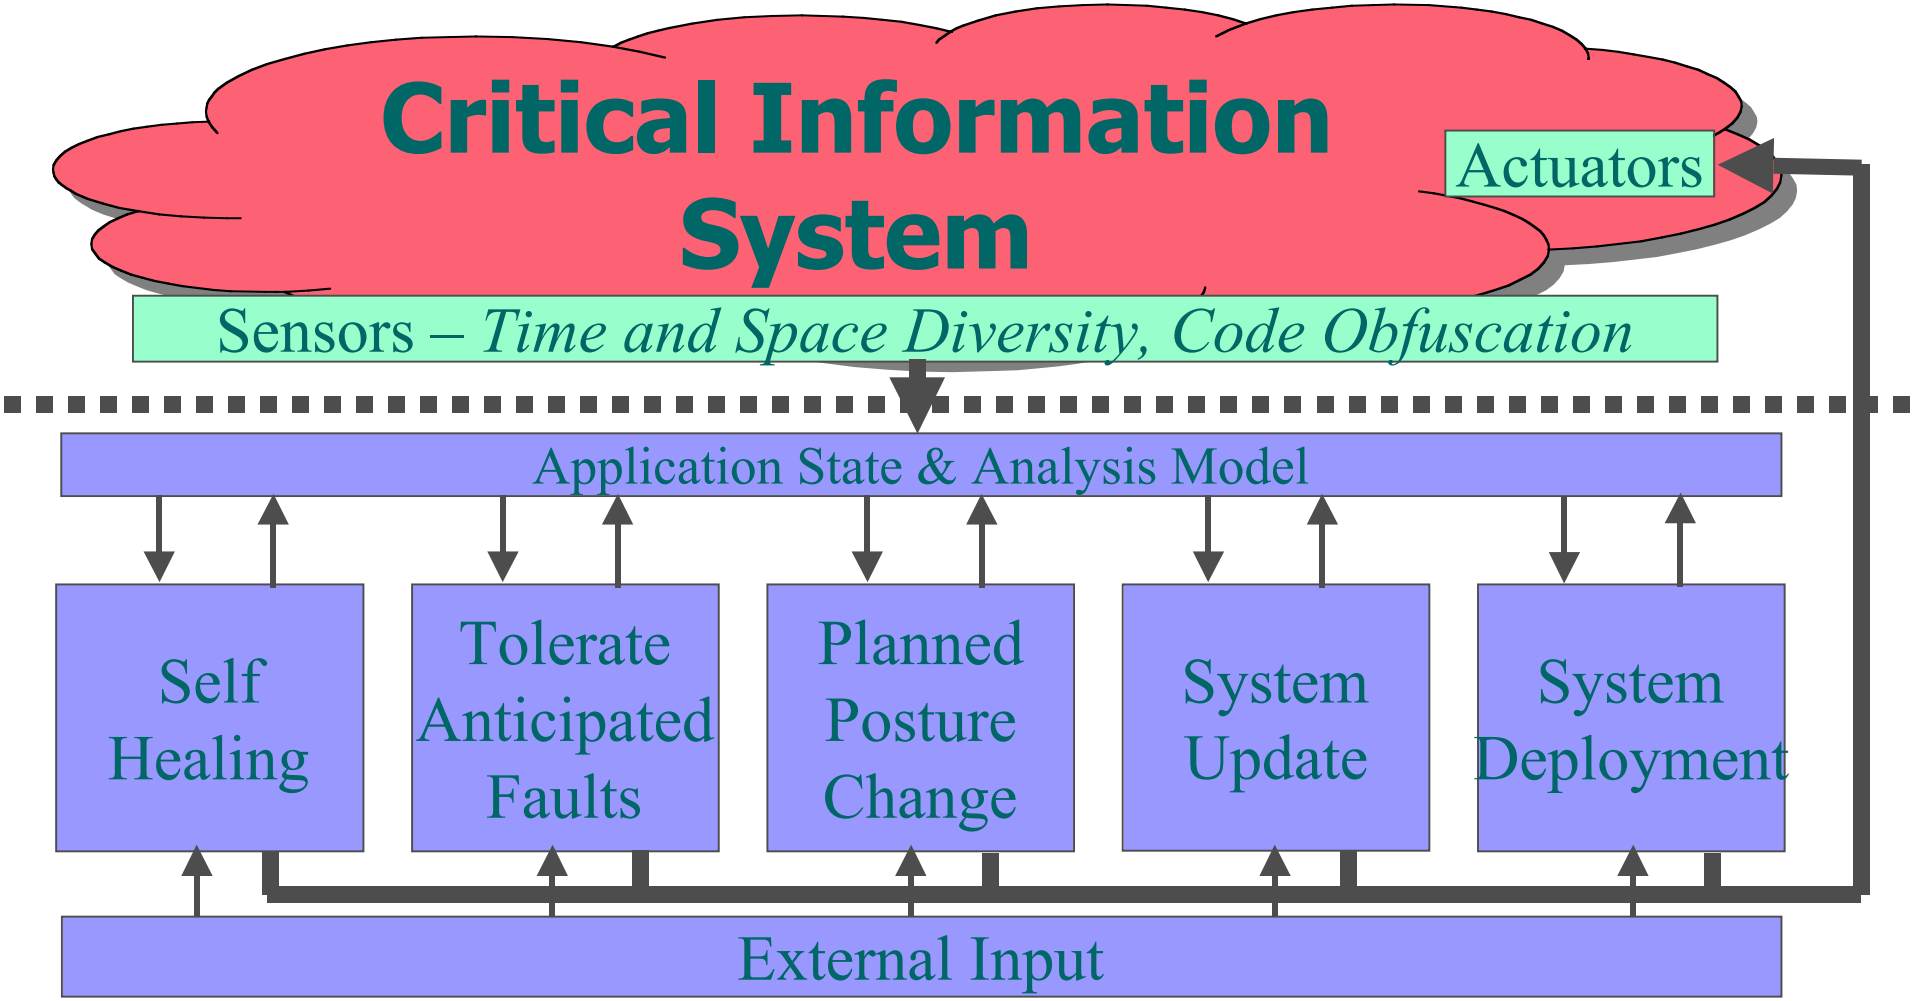

Basic ITUA Architecture

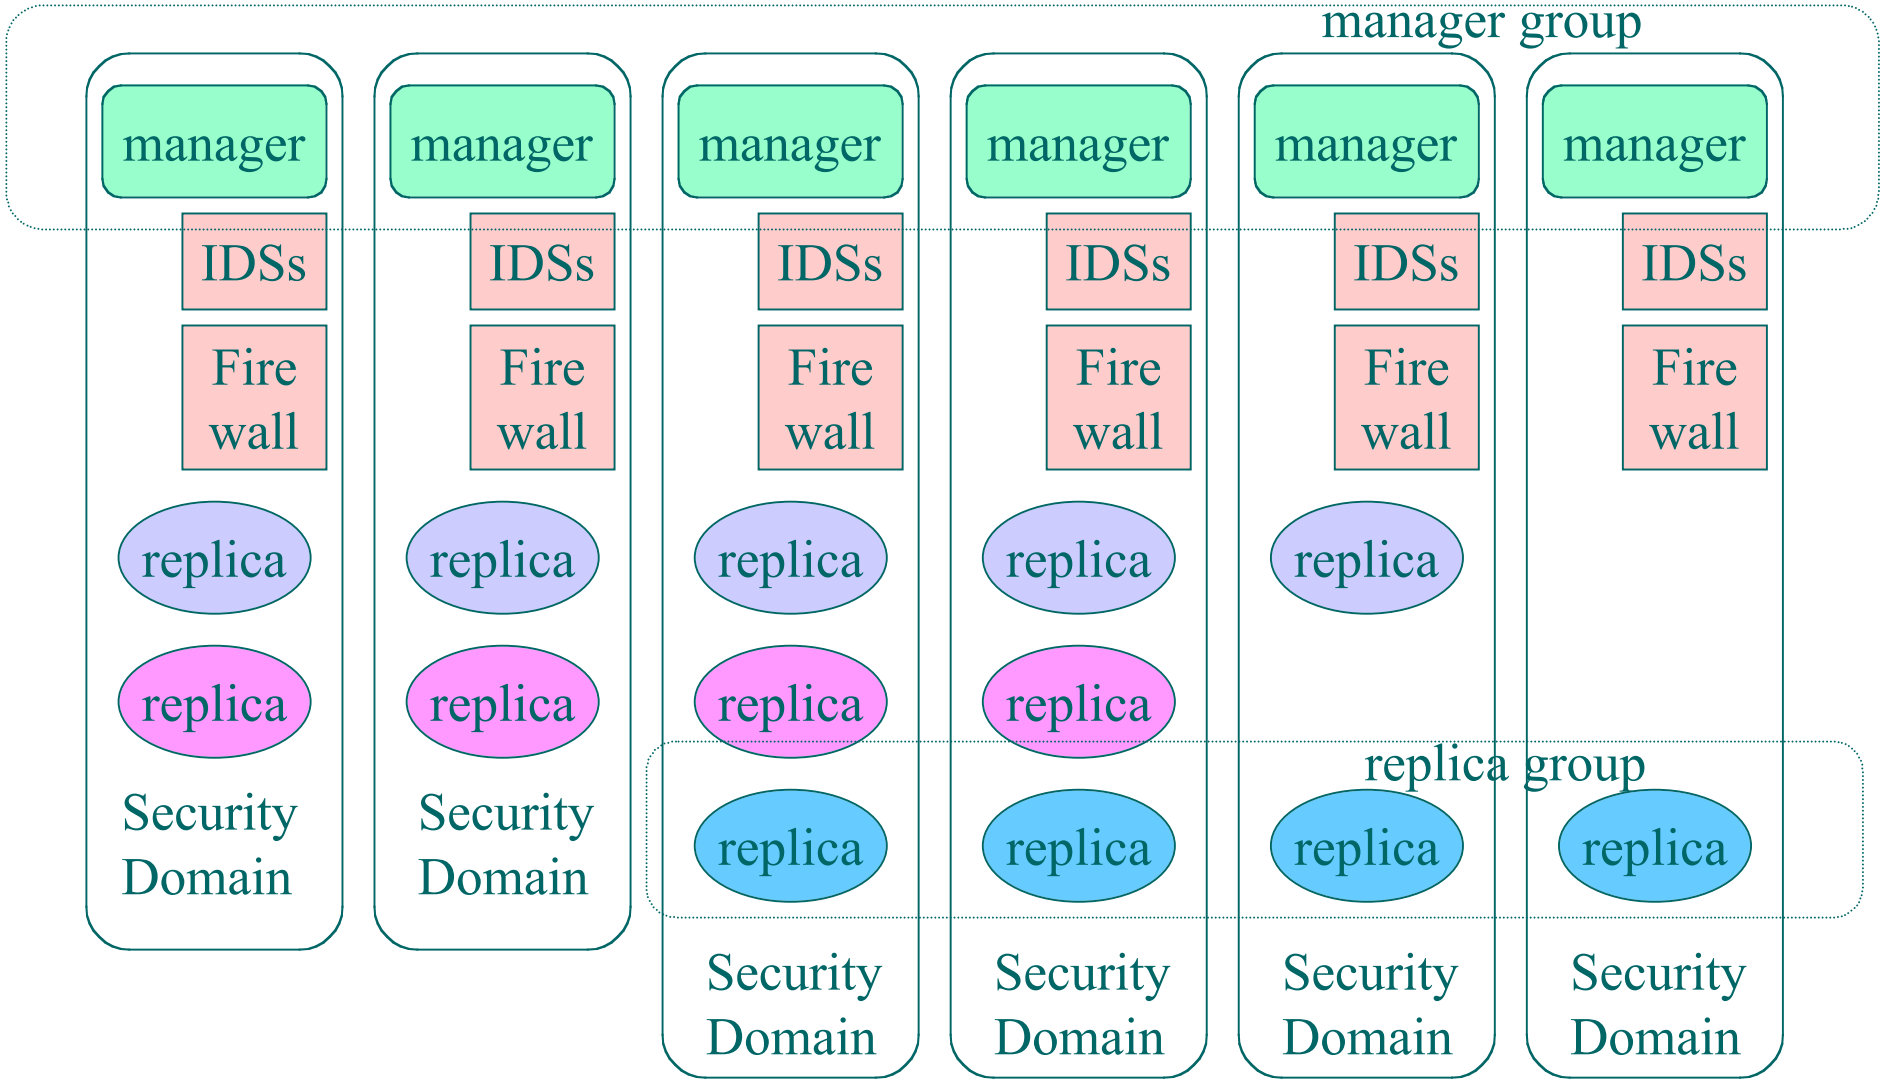

ITDB Components

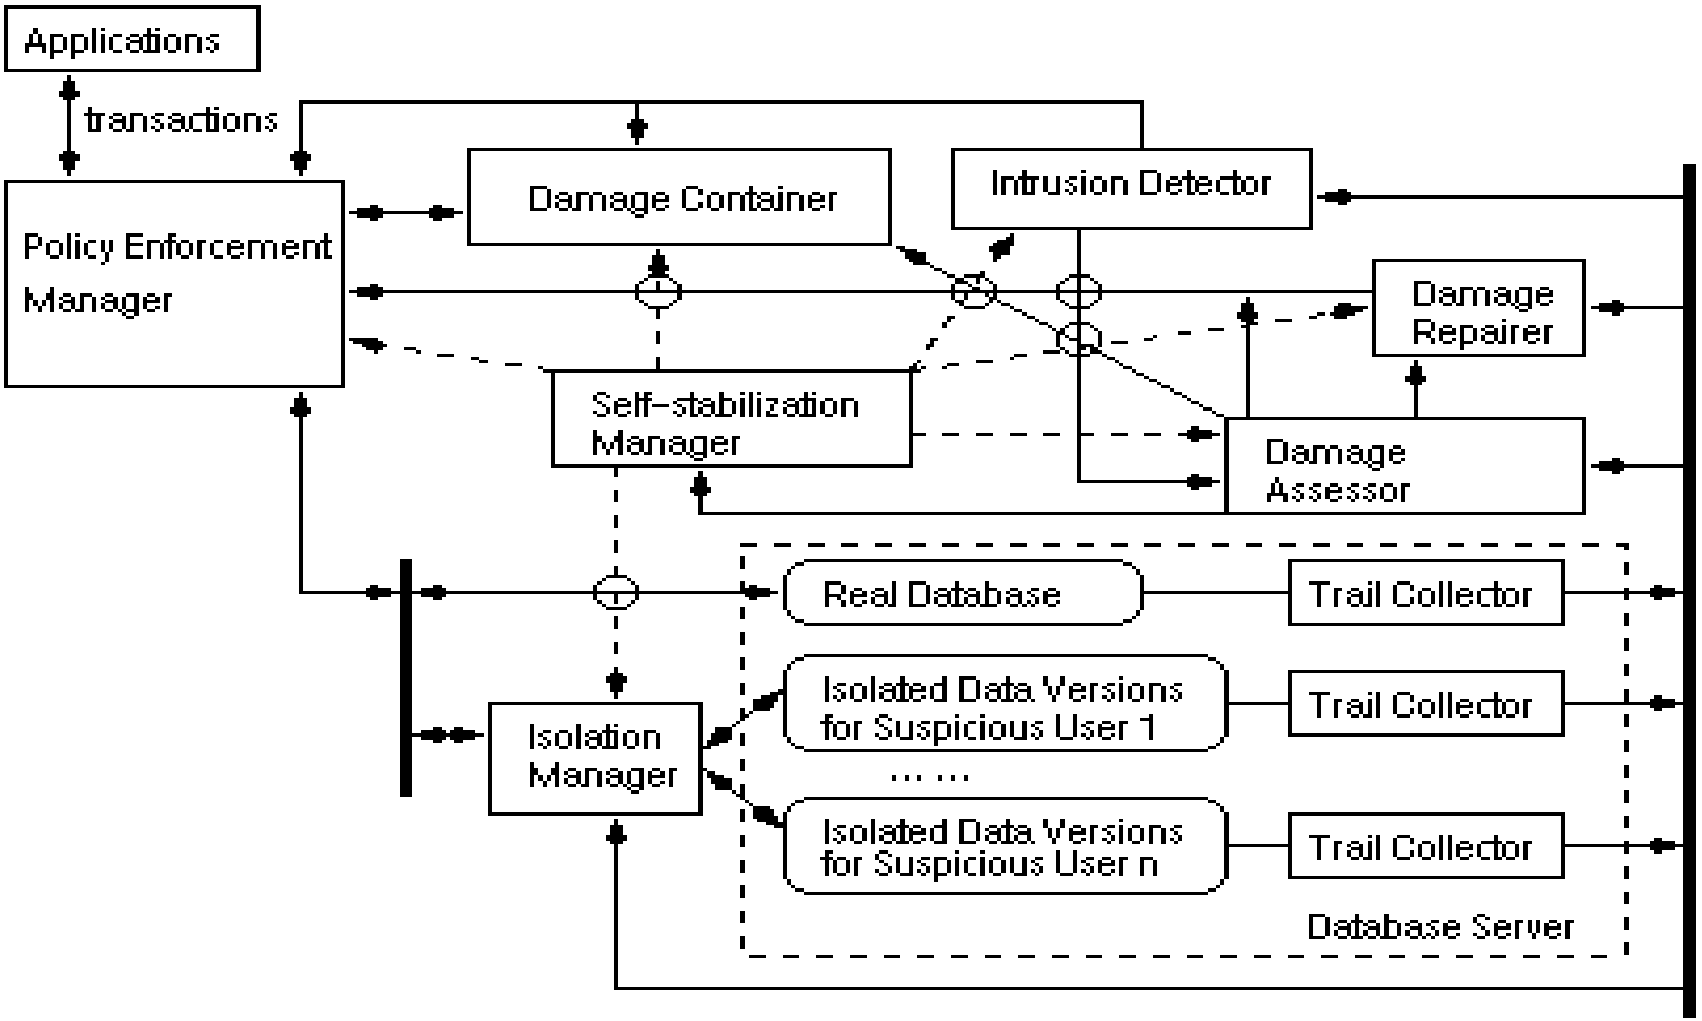

Intrusion Tolerant Middleware

Active Replication on Asynchronous Network

- Each service is replicated over n, n > 1, sites
- Network is intrusion resilient and fault tolerant
  - Bound on message transfer delays is finite but unknown
  - Delay variation does not conform to a known pattern
- Every client request is executed identically on all correct sites
- A majority vote maps sites' responses onto a unique, correct response

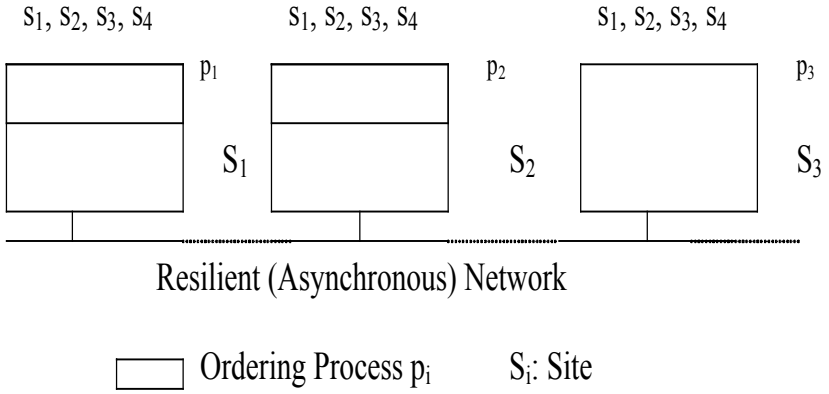

Red Team Panel

Panel Moderator: Steve Bellovin

Panelists:

Fred Avolio, Avolio Consulting, Inc.

Bill Cheswick, Lumeta, Inc.

Sekar Chandersekaran, Institute for Defense Analyses

The red team will address each of the systems, considering the realism of the assumptions, weaknesses in the structures presented, attacks that could defeat the systems' goals.

Each architecture presenter will have an opportunity to respond briefly to the red team's comments.

Audience participation will be invited as time permits

Red Team Panel comments

- Concerns across several systems
  - dependence on good quality intrusion detection, good firewalls
  - uncertainty about diversity (quantifiability, achievability, cost/benefit)
  - vulnerability to DoS attacks (assumed away, in some cases), including forcing computation of crypto signature functions
  - added complexity of IT mechanisms, risk of automated shutdown
  - probabilistic models for attack distributions may be meaningless
  - desire to separate control channels from data channels
- On the other hand:
  - some systems use several mechanisms, not solely IDS's, for detection
  - IT approach (incl. diversity) may incur less time-to-market penalty than high assurance software development
  - many real attacks are scripted; failing to respond at automated speed may doom other systems
  - several possible kinds of diversity - temporal, crypto, spatial
  - randomization, camouflage viewed favorably

Green Team: R&D Directions

Panel Moderator: Carl Landwehr

Panelists:

Fred Schneider, Cornell University

Michael Reiter, Carnegie Mellon University

John D. McLean, Naval Research Laboratory

Paulo Verissimo, Technical University of Lisbon

Richard Hale, U.S. Defense Information Systems Agency

- Each panelist will have a short time to summarize his view of the red team discussion and to suggest research directions motivated by the system and the discussion.
- Following initial comments from the panelists, audience discussion will be invited.

R&D Panel Issues - 1

- How to determine the appropriate assumptions for a particular ITS architecture? (e.g., sync network model may enable DoS attacks)
- How to express the security policy for a particular critical business process / application?
  - Application level security policies are crucial, yet hard to express
  - some IT mechanisms (eg repl) at odds with some rqmnts (conf.)
- How to make better use of architectures employing intrusion masking (vs. detection) such as multiparty computation, bio analogs?
- How to quantify actual diversity of alternative implementations, and benefits thereof?
  - What are the limits of IT approach vs. “high grade” security?
  - How to apply diversity in practice when exact compare may fail?
  - How to deal with relaxation of ACID properties in ITDB and similar architectures?

R&D Issues - 2

- How to model attackers/attacks?
- What is the range of possible responses: shut down, isolate, reboot, but is that all there is (e.g., what about adaptation)?
  - How to express the properties to which adaptive system should converge?
- How to express assurance arguments for ITS; how will they differ from those for safety-critical systems?
- What are the right things to sense (to detect damage / intrusion)?
- How to quantify IT performance vs. cost for different technologies, configurations, architectures?
- How to capture survivability of critical "business process" as a whole?
- How to develop coherent, analyzable intrusion tolerant system architectures?

Other R&D Question

- How can we make progress in this field other than through large scale system demonstration and red-teaming?

Revisiting 1999 pre-OASIS ITS Workshop

- Functions identified as useful for intrusion tolerance
  - *Detection*
  - *Recovery - state restoration*
  - *Masking / error correction*
  - *Redundancy management*
  - *Adaptation / Reconfiguration*
  - *Latent attack detection / self test*
  - *System Behavior Models*
  - *Extent of Compromise - Data Flow Models*
- Potential research directions:
  - *Camouflage - changing protocols to disguise behavior.*
  - *Dynamic Confinement and Authentication.*
  - *Randomness in Algorithms*
  - *Dynamic Reconfiguration and Adaptation*
  - *Fragmentation, Redundancy, and Scattering*
  - *Models and analytical techniques*
  - *Validation/evaluation*
  - *Security Policy for Intrusion Tolerant Systems*
  - *Functional / Analytic Redundancy*
  - *Massive Redundancy.*
  - *Intrusion Tolerant Transaction Processing Schemes*
- Potential Challenge Problems:
  - *Moles on the design team*
  - *Evaluation of ITS*
  - *SW architectures supporting LP or other properties of interest*
  - *Ability to tolerate unanticipated functions / feature interactions / flaws in COTS*
  - *X% of critical functions maintained for n hours following intrusion*

Backup

Elements of Intrusion Tolerance

- Functions
  - intrusion prevention
    - flaw prevention / detection / removal
    - camouflage, obfuscation
  - intrusion detection, correlation, alerting
  - intrusion masking
    - redundancy, voting
  - damage detection
    - redundancy in various forms
    - sensors
  - repair and recovery
  - response
    - resetting defense levels
    - learning new attack signatures

- Technologies
  - quantifying level of intrusion tolerance
    - of a component
    - of an architecture
  - quantifying performance / cost of IT architecture