

# Large-Scale System Security

Steven M. Bellovin

<http://www.cs.columbia.edu/~smb>

Columbia University

(Joint with Sal Stolfo and Angelos Keromytis)

November 13, 2006

Large-Scale Systems

Large-Scale Systems  
Are Different

Scaling Up  
Challenges

Structure

Intrusion Detection

Patching

Conclusion

“Life was simple before World War II.  
After that, we had systems.”

—*Admiral Grace Murray Hopper*

# Large-Scale Systems Are Different

Large-Scale Systems

Large-Scale Systems  
Are Different

Scaling Up  
Challenges

Structure

Intrusion Detection

Patching

Conclusion

- Most interesting *systems* are far more than one single computer
- Rather, they're a collection of many specialized computers
- Sometimes, a single function is replicated, for reliability or load-sharing
- Beyond that, there are many different roles to be filled

# Scaling Up

Large-Scale Systems

Large-Scale Systems  
Are Different

Scaling Up

Challenges

Structure

Intrusion Detection

Patching

Conclusion

- “On the Internet, scale is the only interesting problem” (Mike O’Dell)
- We (more or less) know how to administer a single machine
- We can even make reasonable stabs at groups of identical machines
- The challenge is how to protect large numbers of heterogeneous machines

# Challenges

Large-Scale Systems

Large-Scale Systems  
Are Different

Scaling Up

Challenges

Structure

Intrusion Detection

Patching

Conclusion

- System administration: patches, configuration
- Topology and routing
- Intrusion detection: diverse legitimate input traffic
- Monoculture *and* diversity

Large-Scale Systems

Structure

**System Design**

Diversity versus

Monoculture

System Structure

Intrusion Detection

Patching

Conclusion

- Web server(s)
- Back-end database(s)
- Network Operations Center
- Customer Care
- Backup
- System administration
- Content supplier
- External sources and sinks
- Etc.

# Diversity versus Monoculture

Large-Scale Systems

Structure

System Design

**Diversity versus  
Monoculture**

System Structure

Intrusion Detection

Patching

Conclusion

- If there is too much commonality, a single flaw can take out an entire site
- If there is too much diversity, there may be too little knowledge of system-specific issues
- We probably have the worst of both worlds here: any given function is probably running on the same platform each time

# System Structure

Large-Scale Systems

Structure

System Design

Diversity versus

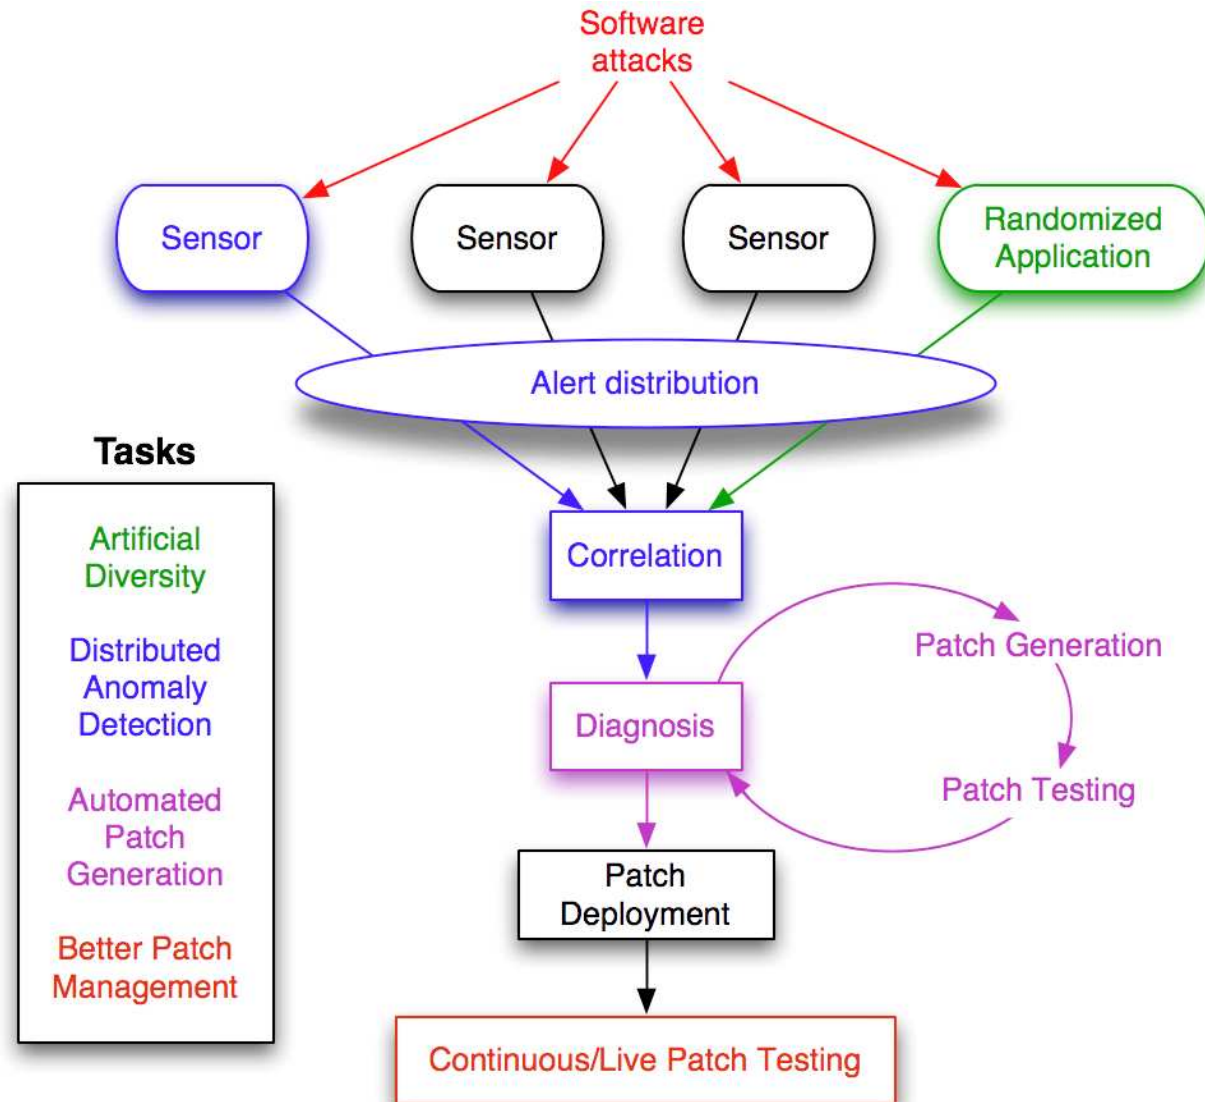
Monoculture

System Structure

Intrusion Detection

Patching

Conclusion



Large-Scale Systems

Structure

**Intrusion Detection**

Intrusion Detection

DNAD

*n*-Gram Detection

Bloom Filters

Patching

Conclusion

# Intrusion Detection

# Intrusion Detection

Large-Scale Systems

Structure

Intrusion Detection

**Intrusion Detection**

DNAD

*n*-Gram Detection

Bloom Filters

Patching

Conclusion

- Can't monitor systems individually
- Instead, take advantage of scale
- Correlate attack from different sensors

Large-Scale Systems

Structure

Intrusion Detection

Intrusion Detection

**DNAD**

*n*-Gram Detection

Bloom Filters

Patching

Conclusion

- Aggregate and compress at different points
- Prioritize alerts
- Use Bloom filter for efficiency and privacy

Large-Scale Systems

Structure

Intrusion Detection

Intrusion Detection  
DNAD

**n-Gram Detection**

Bloom Filters

Patching

Conclusion

- Train on normal data
- (Optionally, train on known-malicious data)
- Compare  $n$ -grams of input to data bases

# Bloom Filters

Large-Scale Systems

Structure

Intrusion Detection

Intrusion Detection

DNAD

*n*-Gram Detection

**Bloom Filters**

Patching

Conclusion

- Again, Bloom filters are used for efficiency and privacy
- Share Bloom filters of suspicious packets among sites
- Highly accurate; detects polymorphic worms

Large-Scale Systems

Structure

Intrusion Detection

**Patching**

Patches and the  
Need for Stability

Automated Patch  
Generation

Automated Patch  
Testing

Artificial Diversity

Other Activities

Conclusion

# Patching

# Patches and the Need for Stability

Large-Scale Systems

Structure

Intrusion Detection

Patching

**Patches and the  
Need for Stability**

Automated Patch

Generation

Automated Patch

Testing

Artificial Diversity

Other Activities

Conclusion

- You can't install patches immediately — there's too much chance they'll break crucial applications
- You can't defer installing patches — most exploits come out after the patch is release
- What about the true "0-day"?

# Automated Patch Generation

Large-Scale Systems

Structure

Intrusion Detection

Patching

Patches and the  
Need for Stability

**Automated Patch  
Generation**

Automated Patch  
Testing

Artificial Diversity

Other Activities

Conclusion

- Detect suspicious traffic
- Analyze and try software patches
- Update production server
- Use source-to-source transformation with templates for known problems

# Automated Patch Testing

Large-Scale Systems

Structure

Intrusion Detection

Patching

Patches and the  
Need for Stability

Automated Patch  
Generation

**Automated Patch  
Testing**

Artificial Diversity

Other Activities

Conclusion

- Dual execution of patched and unpatched versions
- Replicate inputs; compare outputs
- Monitor for correct behavior or signs of intrusion
- When happy, commit the patched version

Large-Scale Systems

Structure

Intrusion Detection

Patching

Patches and the  
Need for Stability  
Automated Patch  
Generation  
Automated Patch  
Testing

**Artificial Diversity**

Other Activities

Conclusion

- Randomize instruction set
- Dual execution
- Detect “part-time failures”
- These indicate suspect code and/or an attack
- Note: this form of diversity does not increase system administration overhead

Large-Scale Systems

Structure

Intrusion Detection

Patching

Patches and the  
Need for Stability

Automated Patch  
Generation

Automated Patch  
Testing

Artificial Diversity

**Other Activities**

Conclusion

- System administration: very little research work in the field
- Human factors of security: large-scale systems deal with very many people
- Infrastructure work: DNS and routing

Large-Scale Systems

Structure

Intrusion Detection

Patching

Conclusion

Conclusion

- We have made scale work for us
- Multiple sensors, aggregation, correlation
- Can handle patches without disrupting operational environment