Red Teaming for Program Managers

SAND2007-4754C

Sandia National Laboratories

Kevin D. Robbins, John F. Clem, Raymond C. Parks, Michael J. Skroch

29 April – 4 May 2007

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

DOE national security laboratory

- FFRDC, GOCO

Our primary mission is nuclear weapons

- Responsible for more than 95% of weapon components

Broader mission in science and engineering to meet national needs

More than 1/4 of our work supports DOD and intelligence community

Sandia National Laboratories

is distributed across many sites

Sandia, New Mexico

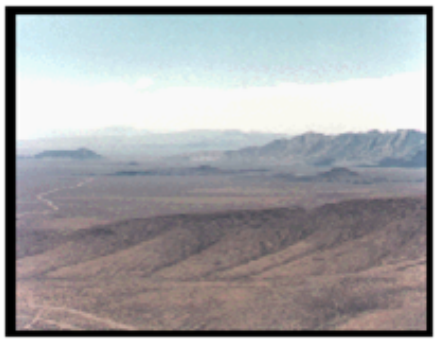

Yucca Mountain, Nevada

WIPP, New Mexico

Tonopah Test Range, Nevada

Kauai Test Facility, Hawaii

Sandia, California

SNL Information Assurance Capabilities

Our mission space means we must protect the information we receive, generate, process, transmit and store.

Defense against advanced threats,

- Trusted foundry,

Advanced research in C4ISR,

- Secure communication protocols (dedicated crypto org),
- Cyber defense tools for analysis and prevention,

Assessment methods development,

Advisors to the military and government, and

Hundreds of technical staff working in IA.

Network Visualization Tools

These visualization tools allow a human analyst to rapidly assess large real-time data sets. These tools are also useful in demonstrating effects of cyber-attacks against SCADA systems.

Host-based Views		Network-based Views
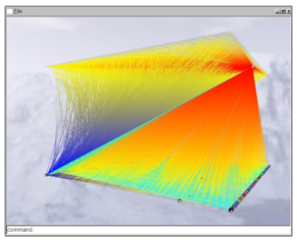	Network under DDoS attack	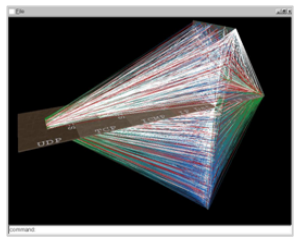
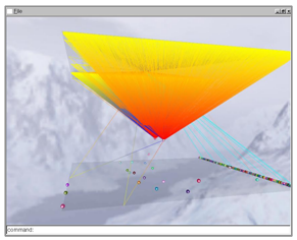	Network port-scanned by NMAP	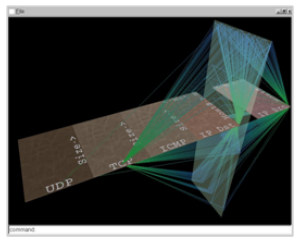
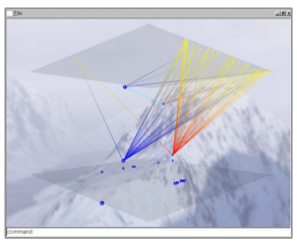	Network probed by NMAP Firewall Probe	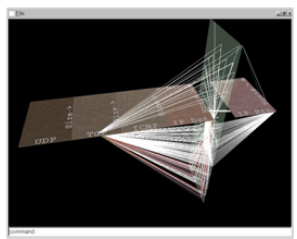

Virtual System Environments

Virtual system environments combine physical and virtual components to simulate and analyze networks and critical infrastructures. These virtual environments support large-scale analysis involving real equipment that is not feasible by other means.

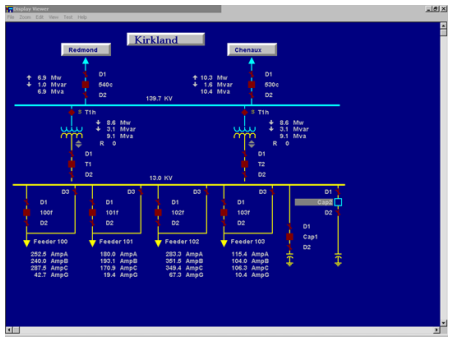

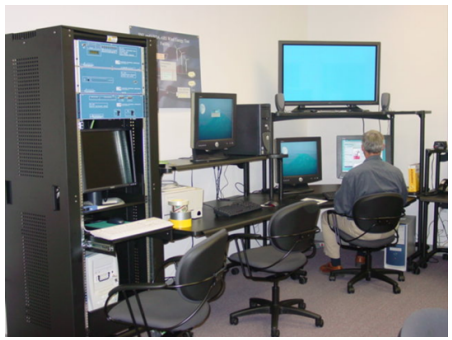

SCADA Linux Appliance

Multipurpose Network Appliance to Improve Infrastructure Security

The SCADA Linux Appliance provides SCADA networks and hosts with many of the security capabilities found in traditional IT elements.

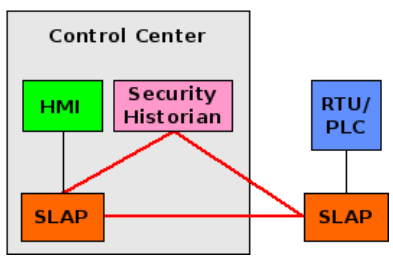

In contrast to some security solutions, each device can be configured to protect single hosts or entire network segments.

The device also provides logging to a centrally-located hardened security historian server for forensic analysis and system visualization.

Multiple Approaches and Resources for Red Teaming

developed to meet the needs of an evolving mission

National security needs have driven Sandia to:

- Expand its set of technical competencies, and
- Develop quality methodologies and techniques for design-assessment processes.

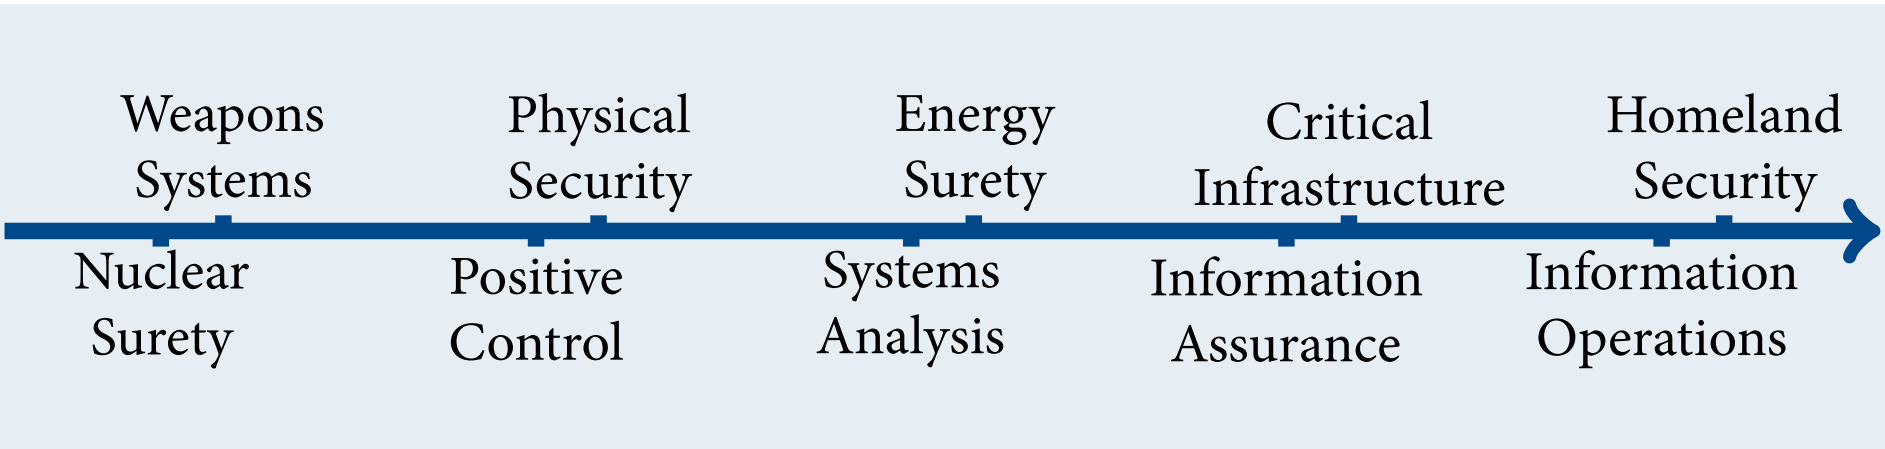

Our Motivation

for providing Red Teaming for Program Managers (RT4PM)

The Information Design Assurance Red Team (IDART) at Sandia National Laboratories has worked

- to improve the process and technique of red teaming.

Our experience has led us to conclude that many of the biggest obstacles to successful assessments have more to do with

- *why* the assessment is needed,
- *what* the red team must deliver,
- *who* performs the assessment, and
- *how* the deliverables will be used to satisfy the assessment goals.

Our goal is to see red teaming become a science and engineering based tool to advance the field of security engineering.

Why Red Teaming for Program Managers?

red teaming addresses the question 'Secure from whom?'

For program managers who must deliver secure components and systems, red teaming is one important means of

- understanding threats and
- exploring effects, impacts, and consequences of adversary actions.

Secure from whom?

Secure components and systems are those that work as intended and only as intended even when an adversary tries to make them do otherwise.

So, it is always important to ask the question, 'Secure from whom and with what knowledge, skills, and tools?'

Why Red Teaming for Program Managers?

the RT4PM process helps meet your objectives

Helps you specify and communicate need for assessment:

- Outline and consider needs to assist interacting with assessors
- Set scope, bounds, constraints to meet budget, timeline, avoid waste
- Start with: objectives, deliverables, available team
- Help define requirements for your team, a BAA, RFQ, Statement of Work (SOW), etc.

Improves your efficiency in using adversary-based assessment

- Leverage existing knowledge in a flexible process you can adapt
- Useful for range of threat, range of lifecycle, physical, cyber, CBRNE, etc.
- Understand how to combine a breadth of options
- Encourages complete consideration of issues (objectives, costs, project impacts, deliverables, assessors)

What is Red Teaming?

a working definition

Red teaming means

- authorized,
- adversary-based
- assessment
- for defensive purposes.

Adversary-based means accounting for

- motivation, goals,
- knowledge, skills,
- tools, and means
- of one or more adversaries.

There are many other types of security assessment that go by various names:

- vulnerability assessment, green teaming, blackhatting, etc.

The Red Teaming for Program Managers Process

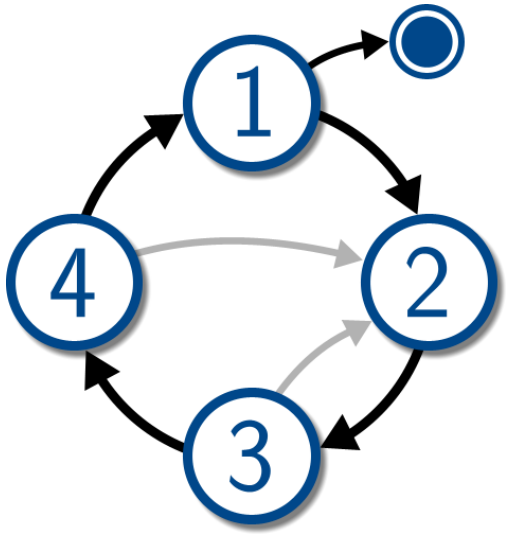

1. Determine your need for red teaming
2. Specify what your red team should do
3. Identify the right red team
4. Plan to use your red team deliverables

When Red Teaming is a Good Choice

Red teaming is a useful tool when one or more of the following are true:

- Hostile, malevolent environment – adversaries,
- Developers more worried about function than security,
- Vulnerabilities are like bugs – where there's one there's more,
- Dynamic, adaptable adversaries,
- Complex systems or system of systems,
- How good/bad is system security,
- New system use that may have unknown consequences,
- Security choices to be made,
- Training and doctrine for the good guys.

When Red Teaming is Not the Best Choice

Red teaming may not be the best choice of security assessment when one or more of the following are true:

- Operational environment is unknown – new or too many,
- Security problems already positively identified,
- Risk or consequence of adversary attack negligible,
- Red team function can be implemented by static model, testbench, or tool,
- Compliance testing or certification is sufficient,
- Not prepared for an extreme answer.

Why Can't Developers do Red Teaming?

independence, objectivity, and knowledge count

Reasons a Program Manager (PM) would prefer an independent, objective red team include:

- You Find What You Look For,
- You Have to Know It When You See It,
- You Find it the Last Place You Look,
- You Can't Say Your Baby is Ugly,
- You Lose Track of the Forest When Cutting Down Trees.

The Red Teaming for Program Managers Process

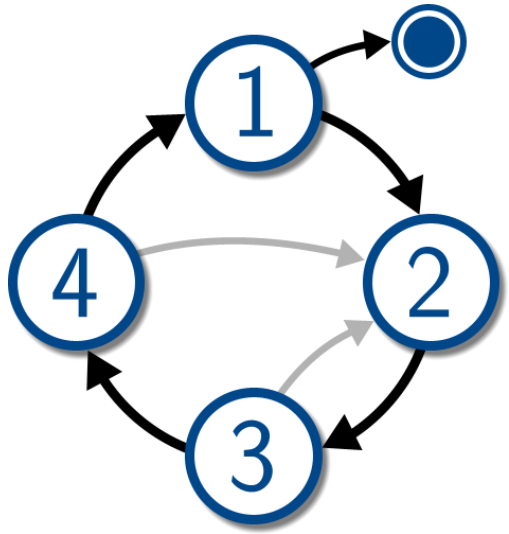

1. Determine your need for red teaming
2. Specify what your red team should do
3. Identify the right red team
4. Plan to use your red team deliverables

Identifying Security Concerns

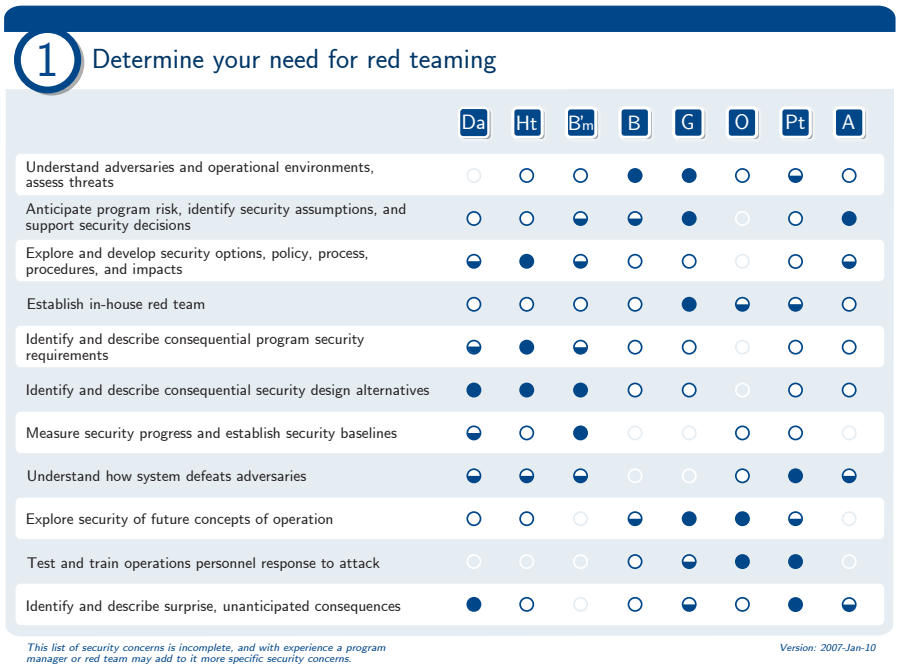

1 Determine your need for red teaming

	Da	Ht	B'm	B	G	O	Pt	A
Understand adversaries and operational environments, assess threats	○	○	○	●	●	○	◒	○
Anticipate program risk, identify security assumptions, and support security decisions	○	○	◒	◒	●	○	○	●
Explore and develop security options, policy, process, procedures, and impacts	◒	●	◒	○	○	○	○	◒
Establish in-house red team	○	○	○	○	●	◒	◒	○
Identify and describe consequential program security requirements	◒	●	◒	○	○	○	○	○
Identify and describe consequential security design alternatives	●	●	●	○	○	○	○	○
Measure security progress and establish security baselines	◒	○	●	○	○	○	○	○
Understand how system defeats adversaries	◒	◒	◒	○	○	○	●	◒
Explore security of future concepts of operation	○	○	○	◒	●	●	◒	○
Test and train operations personnel response to attack	○	○	○	○	◒	●	●	○
Identify and describe surprise, unanticipated consequences	●	○	○	○	◒	○	●	◒

This list of security concerns is incomplete, and with experience a program manager or red team may add to it more specific security concerns.

Version: 2007-Jan-10

Use this table to help determine your need for red teaming.

Security concerns are listed on the left and are

- roughly ordered by project phase,
- from concept to retirement.

Ignore the cross-references on the right.

Your security concern isn't likely to match any of these, but

- Pick one that seems close,
- Pick one that includes yours, or
- Pick one that seems close to your project phase.

Finding Types of Red Team Assessment That May Apply

what types of red teaming could work for your problem

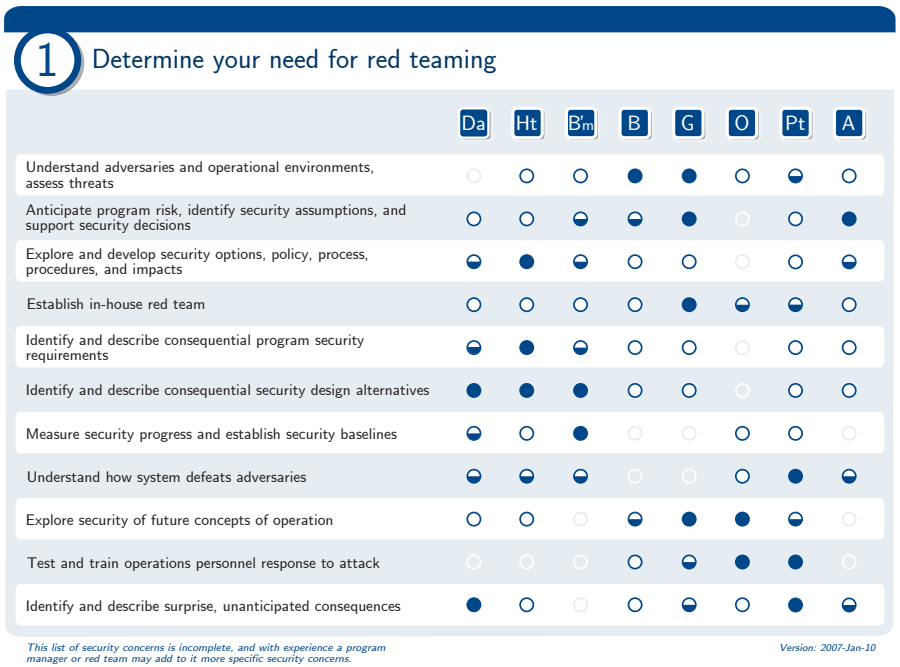

1 Determine your need for red teaming

	Da	Ht	B'm	B	G	O	Pt	A
Understand adversaries and operational environments, assess threats	○	○	○	●	●	○	◒	○
Anticipate program risk, identify security assumptions, and support security decisions	○	○	◒	◒	●	○	○	●
Explore and develop security options, policy, process, procedures, and impacts	◒	●	◒	○	○	○	○	◒
Establish in-house red team	○	○	○	○	●	◒	◒	○
Identify and describe consequential program security requirements	◒	●	◒	○	○	○	○	○
Identify and describe consequential security design alternatives	●	●	●	○	○	○	○	○
Measure security progress and establish security baselines	◒	○	●	○	○	○	○	○
Understand how system defeats adversaries	◒	◒	◒	○	○	○	●	◒
Explore security of future concepts of operation	○	○	○	◒	●	●	◒	○
Test and train operations personnel response to attack	○	○	○	○	◒	●	●	○
Identify and describe surprise, unanticipated consequences	●	○	○	○	◒	○	●	◒

This list of security concerns is incomplete, and with experience a program manager or red team may add to it more specific security concerns.

Version: 2007-Jan-10

Now use the cross-references to find

- the types of red teaming that will provide the assessment structure,
- and the types that may contribute concepts.

Make a list of all types that may apply

- you will decide which to use in the next step.

The Red Teaming for Program Managers Process

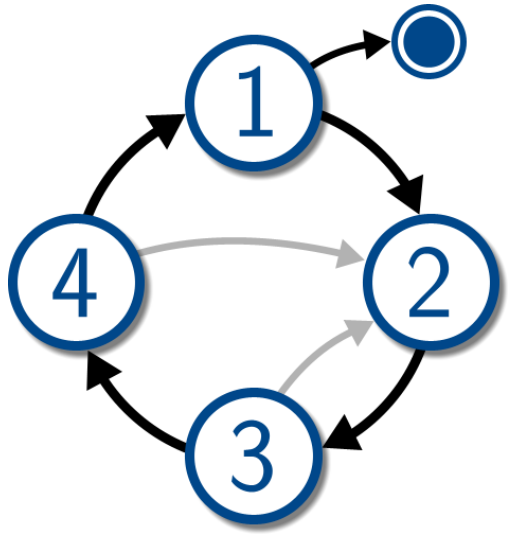

1. Determine your need for red teaming
2. Specify what your red team should do
3. Identify the right red team
4. Plan to use your red team deliverables

Considering Types of Red Teaming That May Apply

You now have a list of types of red teaming that may apply.

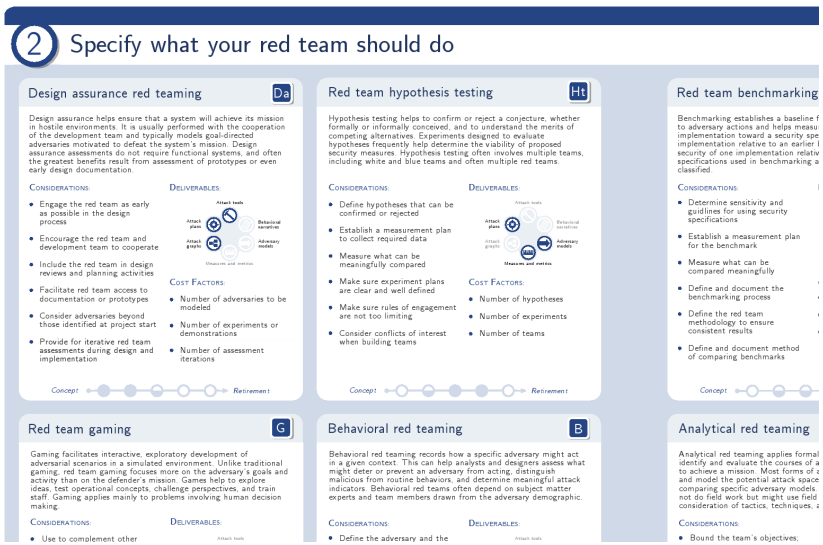

Use this section to find those that do apply

- whether directly, or
- in combination with another type.

Many if not most real world assessments

- require a hybrid approach that blends
- concepts and methods from several types of red teaming.

Use the types of red teaming that do apply to your problem to determine the

- scope and statement of work for the assessment.

Using the Types of Red Teaming

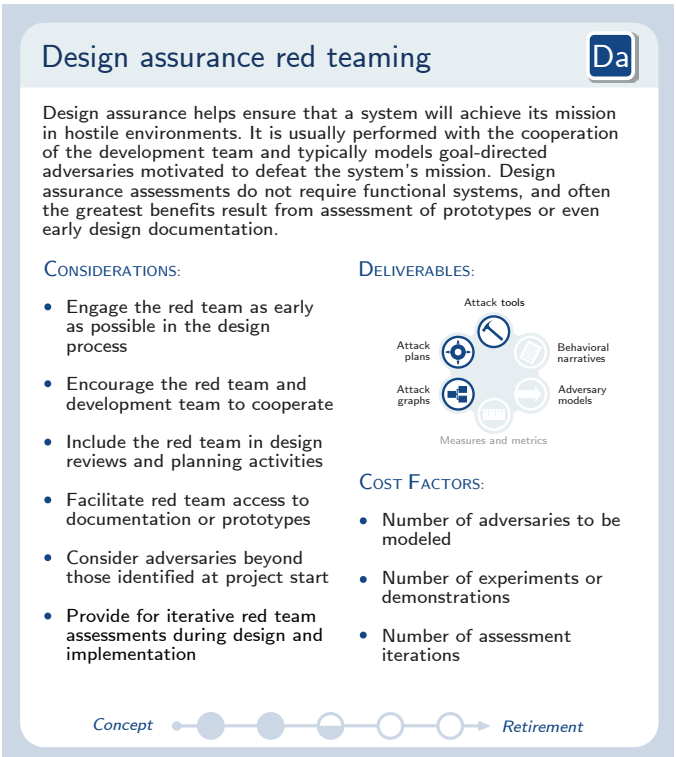

Each type of red teaming, for example *Design assurance*, is identified by

- a description of why it would be used,
- considerations, the do's and don't's,
- deliverables a PM might need,
- factors that drive the assessment cost,
- a suggestion of when to use it.

Design Assurance Red Teaming

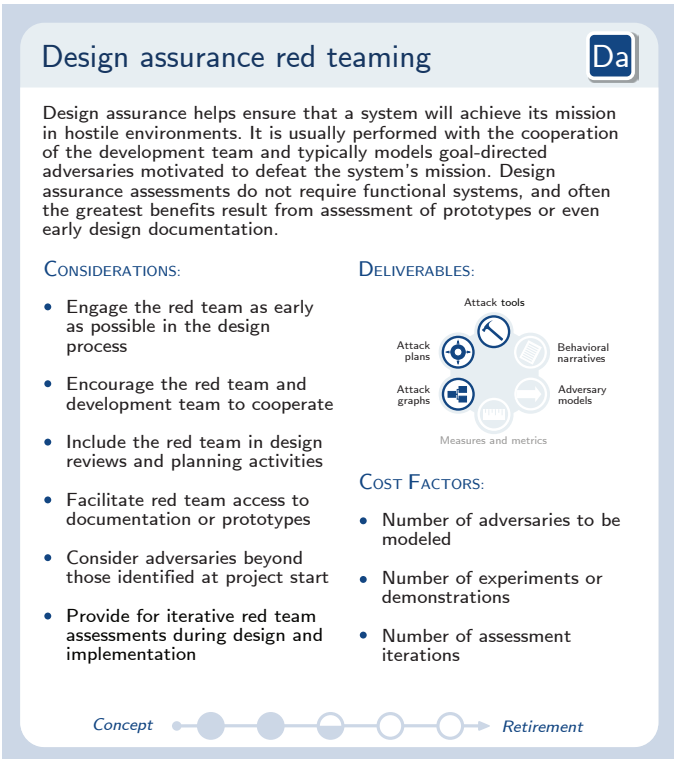

Design assurance red teaming

- helps ensure a system design will achieve its mission in a hostile environment,
- in cooperation with the system developers,
- accounting for adversaries whose goal is to defeat the system's mission.

Red Team Hypothesis Testing

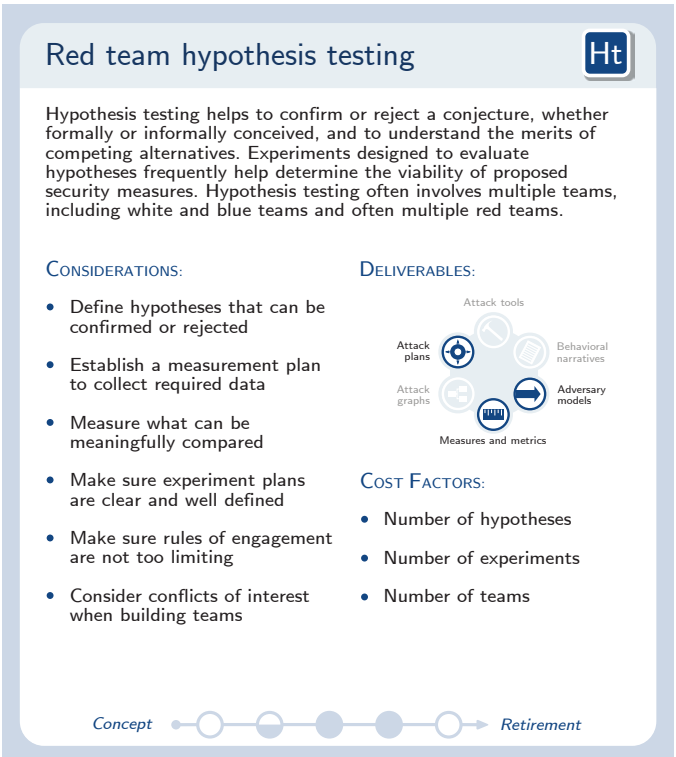

Red team hypothesis testing

- evaluates hypotheses such as *layered, partial defenses are better than a single strong defense* and
- IR *sensors detect adversaries better than radar*, in
- experiments with a blue team, white team, and one or more red teams.

Red Team Gaming

Red team gaming

- supports interactive development of adversarial scenarios,
- focuses on adversary goals and not defender's mission, and
- focuses on human decision making and not specific methods.

Behavioral Red Teaming

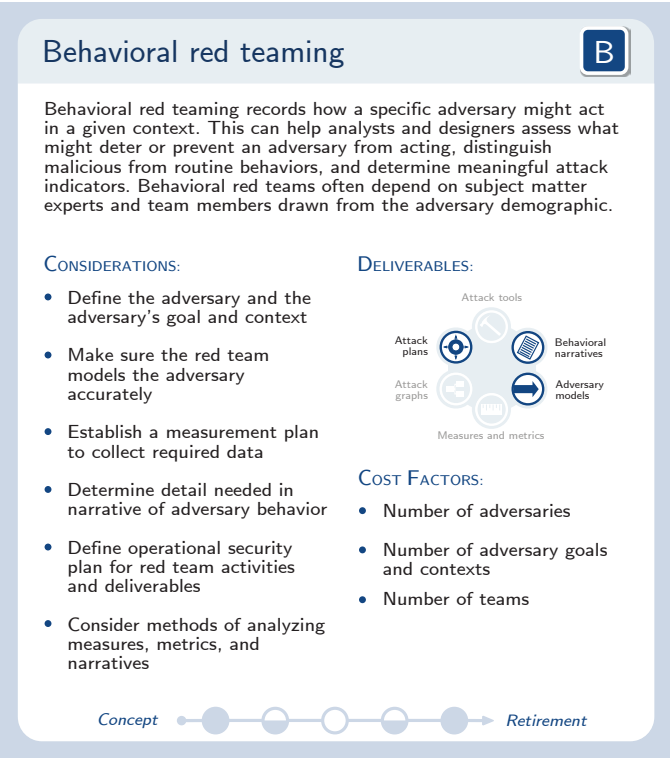

Behavioral red teaming

- uncovers specific adversary actions in a given context
- for indications, warnings, detection, and deterrence and
- relies on subject matter experts often from the adversary demographic.

Red Team Benchmarking

Red team benchmarking

Benchmarking establishes a baseline for comparing system responses to adversary actions and helps measure progress of an implementation toward a security specification, progress of an implementation relative to an earlier benchmark, and measured security of one implementation relative to another. Security specifications used in benchmarking are often sensitive or even classified.

CONSIDERATIONS:

- Determine sensitivity and guidlines for using security specifications
- Establish a measurement plan for the benchmark
- Measure what can be compared meaningfully
- Define and document the benchmarking process
- Define the red team methodology to ensure consistent results
- Define and document method of comparing benchmarks

DELIVERABLES:

COST FACTORS:

- Number of benchmarks
- Number of red teams
- Number of assessment iterations

Concept — *Retirement*

Red team benchmarking

- establishes baselines for system performance under attack,
- measures progress of a system toward its security specification, and
- supports comparison between different implementations.

Operational Red Teaming

Operational red teaming

Operational red teaming models an active adversary within a live or simulated context. Operational red teams seek to defeat the target system's mission in realistic deployment environments. Operational red teaming helps to train staff, conduct testing and evaluation, validate concepts of operation, and identify vulnerabilities. Operational red teams will usually have less time than real-world adversaries to prepare.

Considerations:

- Match effort's fidelity with program requirements
- Carefully weigh and define rules of engagement
- Validate that the red team acts within required constraints
- Obtain all required authorizations; identify and mitigate activity risks
- Consider establishing a standing operational red team or opposing force (OPFOR)
- Match existing scenarios with the program's concept of operations (CONOPS)

Deliverables:

Cost Factors:

- Scope, scale, and complexity of target system
- Materials required to execute program
- Level of active response from blue team

Concept — *Retirement*

Operational red teaming

- simulates an active adversary seeking to defeat the defender's mission
- in a live or simulated, realistic deployment
- to validate CONOPS, identify vulnerabilities, or for training and OT&E.

Analytical Red Teaming

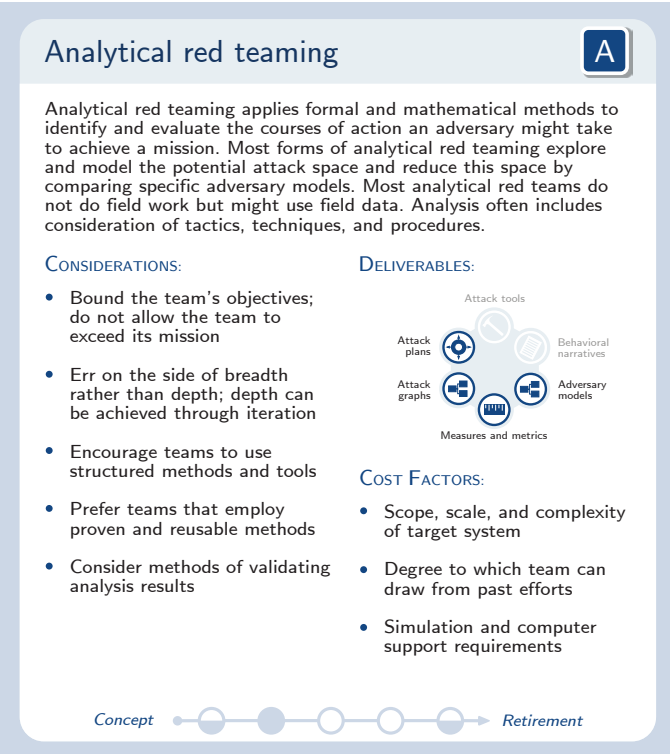

Analytical red teaming

- uses formal and mathematical methods
- to explore attack or consequence space,
- and often considers adversary tactics, techniques, and procedures.

Analytical red teaming can incorporate other types of red teaming to

- mathematically predict the probability of particular (undesirable) consequences.

Penetration Testing

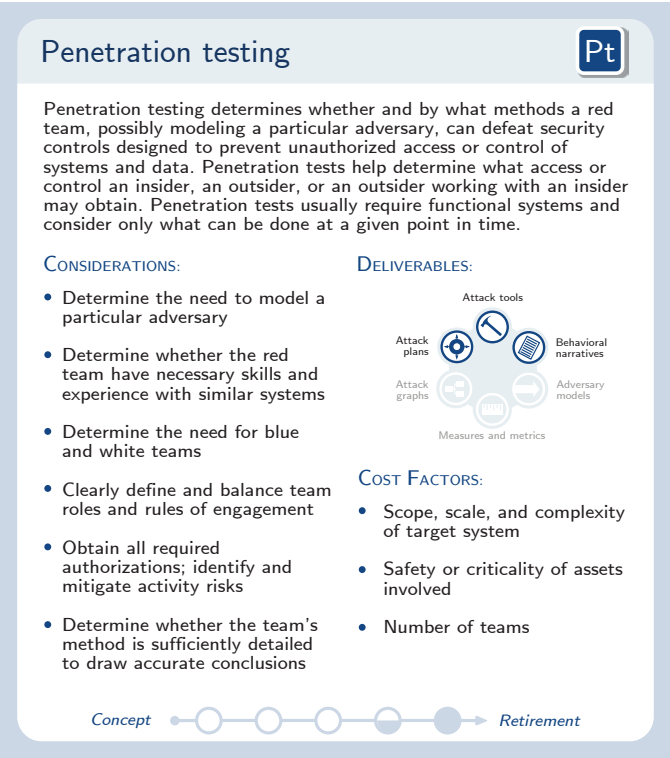

Penetration testing

- involves active determination of methods and tools an adversary needs to attack
- a live, or at least, functional system, and
- measures interaction between the adversary and the system.

Crafting the Scope and Statement of Work

You now have a list of types of red teaming that do apply to your problem.

Use descriptions and considerations of the types that apply to

- identify tasks,
- manage project interdependencies, and
- avoid common problems.

Use deliverables to consider what the red team must provide.

Use cost factors to bound the assessment scope.

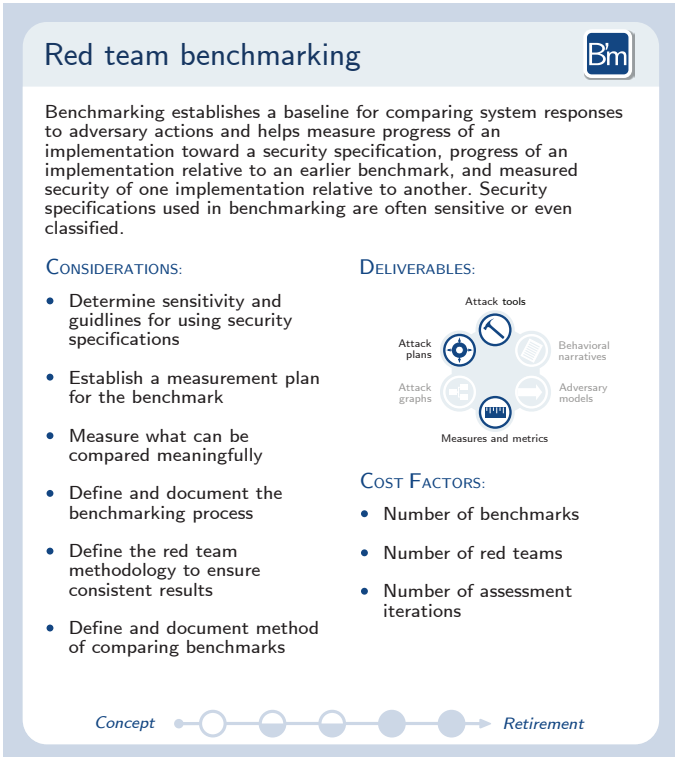

The Red Teaming for Program Managers Process

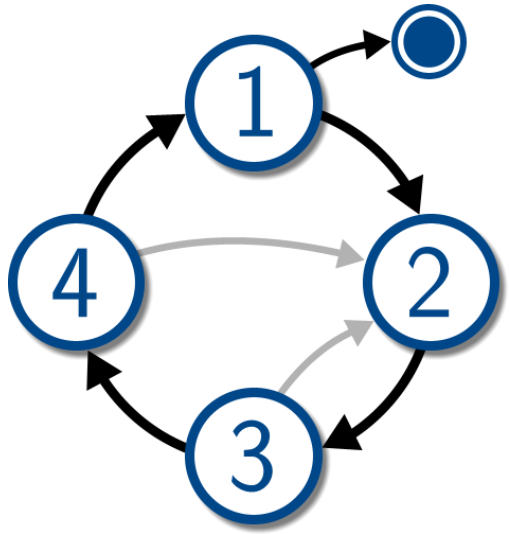

1. Determine your need for red teaming
2. Specify what your red team should do
3. Identify the right red team
4. Plan to use your red team deliverables

Select Important Criteria

that match the scope, statement of work, and program

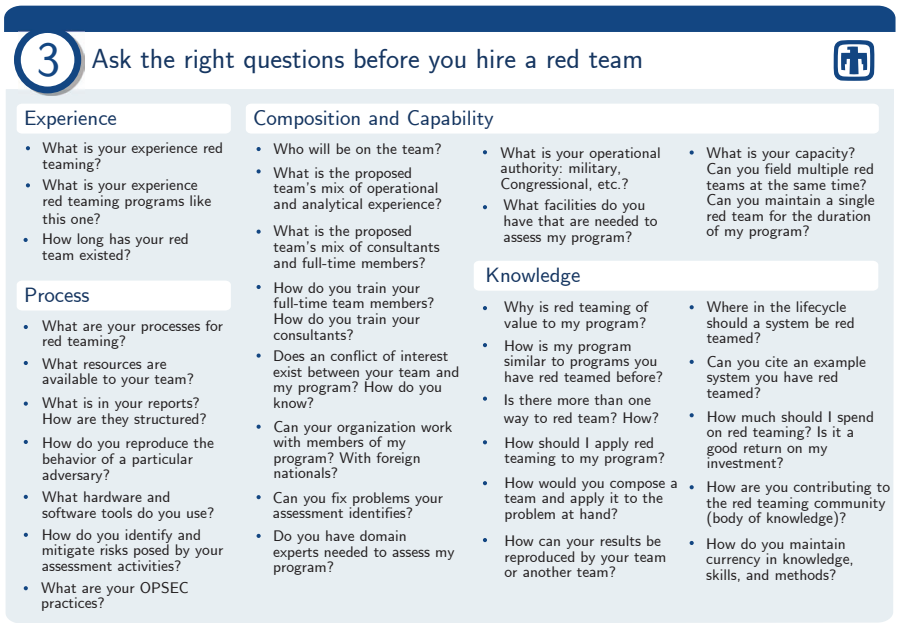

3 Ask the right questions before you hire a red team

Experience

- What is your experience red teaming?
- What is your experience red teaming programs like this one?
- How long has your red team existed?

Process

- What are your processes for red teaming?
- What resources are available to your team?
- What is in your reports? How are they structured?
- How do you reproduce the behavior of a particular adversary?
- What hardware and software tools do you use?
- How do you identify and mitigate risks posed by your assessment activities?
- What are your OPSEC practices?

Composition and Capability

- Who will be on the team?
- What is the proposed team's mix of operational and analytical experience?
- What is the proposed team's mix of consultants and full-time members?
- How do you train your full-time team members? How do you train your consultants?
- Does an conflict of interest exist between your team and my program? How do you know?
- Can your organization work with members of my program? With foreign nationals?
- Can you fix problems your assessment identifies?
- Do you have domain experts needed to assess my program?
- What is your operational authority: military, Congressional, etc.?
- What facilities do you have that are needed to assess my program?
- What is your capacity? Can you field multiple red teams at the same time? Can you maintain a single red team for the duration of my program?

Knowledge

- Why is red teaming of value to my program?
- How is my program similar to programs you have red teamed before?
- Is there more than one way to red team? How?
- How should I apply red teaming to my program?
- How would you compose a team and apply it to the problem at hand?
- How can your results be reproduced by your team or another team?
- Where in the lifecycle should a system be red teamed?
- Can you cite an example system you have red teamed?
- How much should I spend on red teaming? Is it a good return on my investment?
- How are you contributing to the red teaming community (body of knowledge)?
- How do you maintain currency in knowledge, skills, and methods?

Select important criteria that match the scope, statement of work, and program:

- Experience,
- Process,
- Composition,
- Capability, and
- Knowledge.

Recommendations

- Sketch a rough SOW first, using the RT4PM process,
- then develop a list of *must have* and *would like* criteria.

The Red Teaming for Program Managers Process

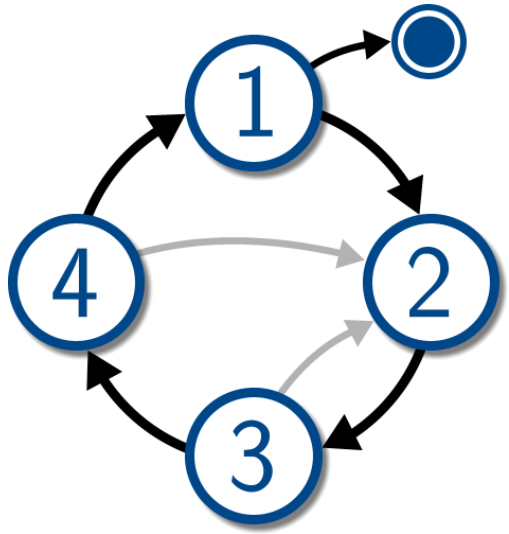

1. Determine your need for red teaming
2. Specify what your red team should do
3. Identify the right red team
4. Plan to use your red team deliverables

Deliverables Close the Loop

Knowing what deliverables are truly needed and what purpose they will serve

- allows a PM to get the best value from red team assessment.

Knowing how to use common red team deliverables will help a PM, analyst, or decision maker

- Determine what deliverables are needed,
- Identify red teams that can provide the needed deliverables, and
- Use the deliverables to meet the program need.

Attack Graphs and Trees

Attack graphs and trees:

- Document attack concepts with high-level metrics for each attack step,
- Develop common understanding of attacks and how they work,
- Explain attack spaces and choices or decisions the adversary will face.

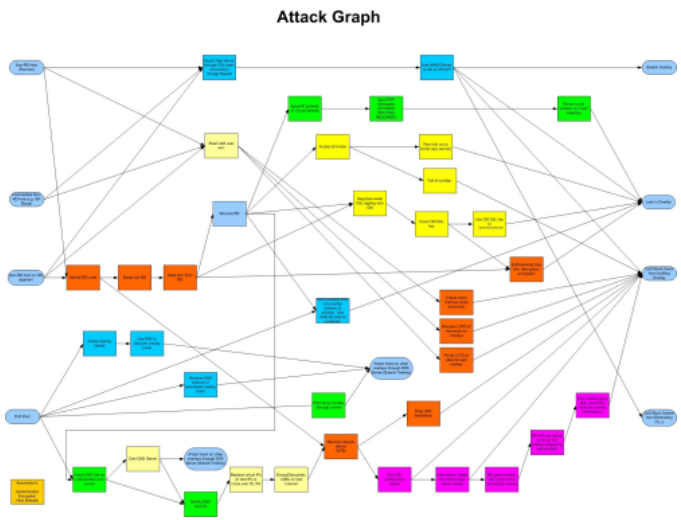

Attack trees usually focus on consequence or vulnerabilty.

Attack graphs usually focus on adversary activity and interaction with the target.

Adversary Models

Adversary models

- define and document the behavior and decisions of an adversary
- in a given context or environment.

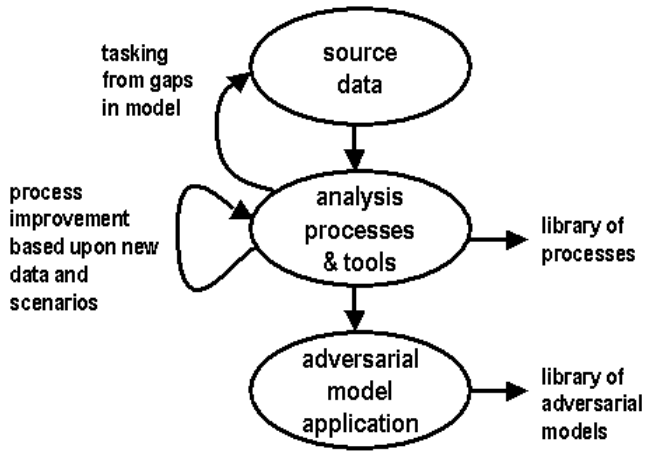

Adversary models vary in complextity from the simplistic to the systematic.

Adversary models:

- Bound red team behavior to that expected of a given adversary,
- Document objectives or goals for the red team to pursue that are consistent with those of a given adversary, and
- Enable more measurable, consistent, and reproducible results.

Considerations for Using Red Team Deliverables

It's important to schedule time to use red team deliverables

- and account for other schedule impacts from the assessment.

For example, staff may need training to use the deliverables.

What quality of deliverables is needed?

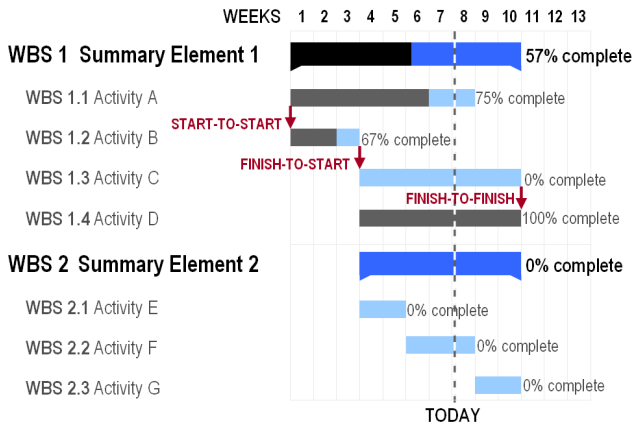

Common Topics in Red Team Assessment

There are a number of common considerations for using red team assessment:

- Depth and breadth,
- Threat assessment,
- Adversary models and scenarios,
- Organizational Conflicts of Interest,
- Access to information,
- Notice,
- Cooperation, and
- Assessment security.

Threat Assessment

Threat assessment and red team assessment are frequently interdependent

- threat assessment can identify adversaries of concern for red teaming, and
- red teaming may be used to enhance threat assessment.

Threat assessment involves

- determining adversary objectives and likely actions,
- identifying needed environments and those who might take action,
- measuring necessary adversary resources in terms of
  - knowledge, skills, tools, and numbers, and
- correlating observed events to likely adversaries.

Notice

Whether or not notice is given, permissions and safeguards are needed in active red team engagements.

No notice efforts

- do not influence target behavior prior to assessment,
- may more realistically model adversaries and their limitations,
- and may allow unbiased retesting.

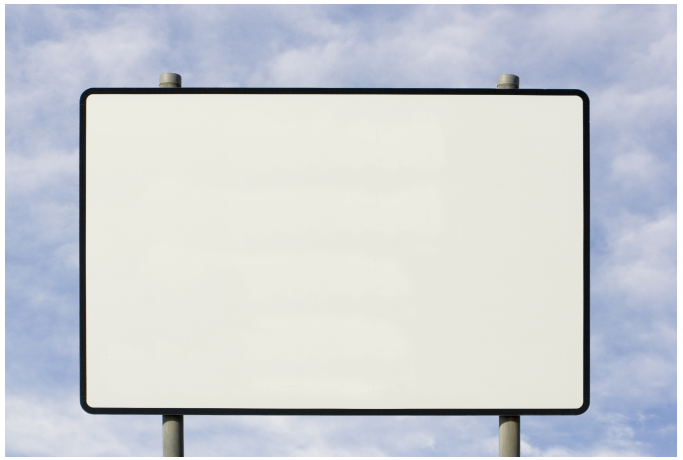

Notice given efforts

- may be useful for an initial engagement prior to a no notice effort,
- may require trust building by the red team, and
- should account for and possibly measure the affect of the notice.

http://www.sandia.gov/redteam2007

- Designed to bring together professional red teamers and the infrastructure protection community
- Unclassified (AUSCANZUKUS)
- Three days
  - First two days: Presentations and workshops
  - Third day: Sandia's Red Teaming for Program Managers (RT4PM) course
- Deadline for presentation proposals: 25 May 2007

http://www.sandia.gov/redteam2007

- Topics and tracks
  - Methods, techniques, and tools
  - Program and project management
  - Information sharing within organizations and across sectors
- Who should attend?
  - Military, Federal, state, and local officials responsible for homeland security, infrastructure protection, and law enforcement
  - Military and intelligence organizations responsible for red teaming and adversary modeling
  - Commercial providers of red teaming services
  - Commercial infrastructure providers

Questions and Open Discussion

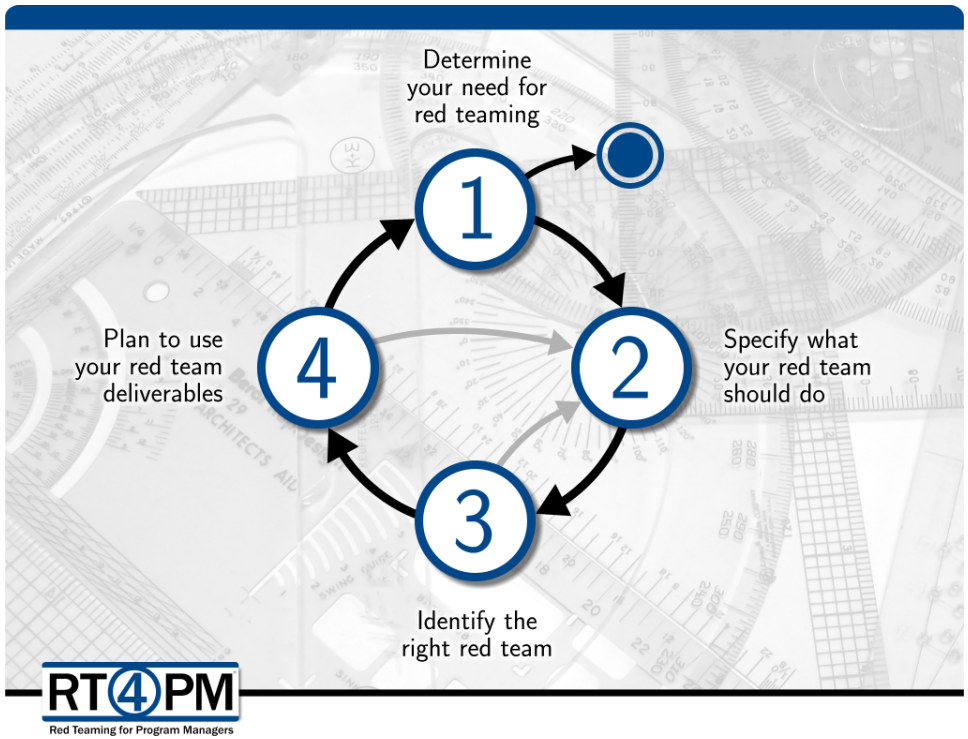

Additional Materials

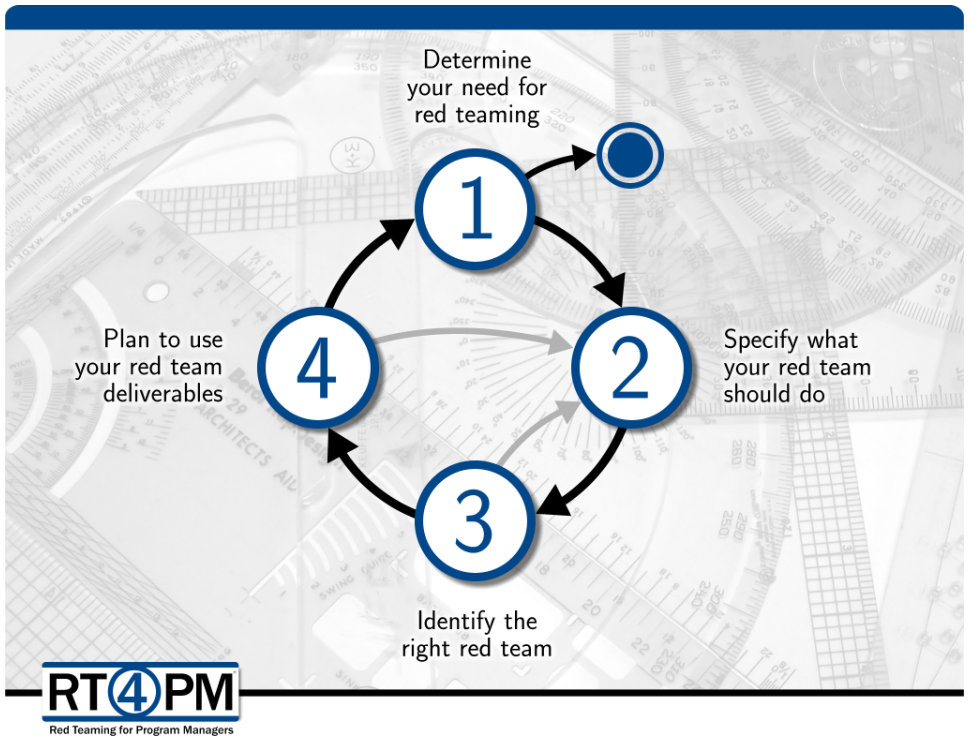

Common Topics in Red Team Assessment

There are a number of common considerations for using red team assessment:

- Depth and breadth,
- Threat assessment,
- Adversary models and scenarios,
- Organizational Conflicts of Interest,
- Access to information,
- Notice,
- Cooperation, and
- Assessment security.

Depth and Breadth

Depth and breadth influence cost, schedule, and team composition, where

- breadth is the diversity of issues, components, geometry, etc. and
- depth is the detail or time applied in each area.

Depth and breadth are important in scoping red team assessment, because

- excessive breadth and depth may waste resources and
- threaten assessment completeness.

Recommendations

- Set flexible bounds on the assessment parameters.
- Expand boundaries upon discovery of issues.

Threat Assessment

Threat assessment and red team assessment are frequently interdependent

- threat assessment can identify adversaries of concern for red teaming, and
- red teaming may be used to enhance threat assessment.

Threat assessment involves

- determining adversary objectives and likely actions,
- identifying needed environments and those who might take action,
- measuring necessary adversary resources in terms of
  - knowledge, skills, tools, and numbers, and
- correlating observed events to likely adversaries.

Adversary Models and Scenarios

Adversary models bound red team behavior and help ensure the red team analysis is sound.

Adversary or threat models may specify

- specific people or groups,
- people from a society, country, or region,
- people with particular technical background and experience,
- people with particular motivation and intent, etc.

Scenarios explain adversary models in realistic context.

Recommendations

- Consider threat as part of your system's environment.
- Consider bounds on adversary and red team behavior.
- Consider a combination or range of threats.

Organizational Conflicts of Interest

Organizational Conflicts of Interest (OCI) may impact your use of a red team by

- actually influencing a red team's results, or
- affecting how red team results are interpreted and valued.

Recommendations

- Consider whether OCI might affect your effort.
  - Will results actually be biased?
  - Who else will use the results?
  - Can people be adversely affected by the assessment?
- Institute controls to mitigate any OCI you identify.

Access to Information

The amount of information given to a red team is a control on cost and process.

Free access to data

- generally reduces assessment cost,
- allows the red team to more easily model sophisticated and complex adversaries, but
- requires more strict process on red team methods.

Little access to data

- generally increases assessment cost,
- may more realistically model adversaries and their limitations,
- and may be useful in some exercises or demonstrations.

Notice

Whether or not notice is given, permissions and safeguards are needed in active red team engagements.

No notice efforts

- do not influence target behavior prior to assessment,
- may more realistically model adversaries and their limitations,
- and may allow unbiased retesting.

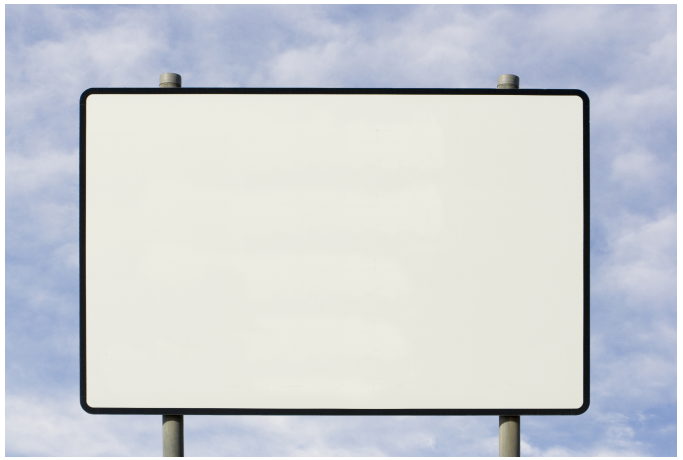

Notice given efforts

- may be useful for an initial engagement prior to a no notice effort,
- may require trust building by the red team, and
- should account for and possibly measure the affect of the notice.

Cooperation

Cooperation involves sharing of information, access to systems, and participation by the target.

Red team assessments without cooperation

- may cost more,
- usually require a white team or go-between,
- and usually require the red team to gather information from open source or other means.

Cooperative red team assessments

- usually cost less,
- often uncover more information and vulnerabilities,
- are more likely to result in immediate improvements, and
- requires trust building by the red team.

Assessment Security

Vulnerabilities and other issues the red team discovers may be sensitive or even classified.

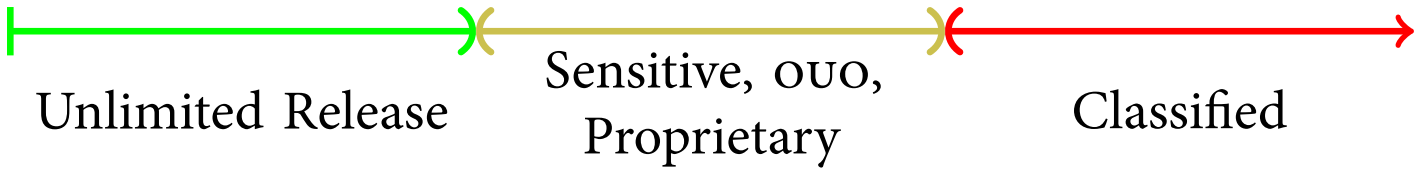

Unlimited release

- often relates to systems that are not operational or components not used in operational systems,
- is often found in R&D environments.

Sensitive, OUO, proprietary

- may have export control implications, and
- requires OPSEC controls and other security measures.