

# EP2790 Security Analysis of Large-Scale Computer Systems

## Motivation

Companies today have thousands of software based computer systems that all are depending on one another in a large complex network, a system-of-systems. That IT attacks succeed to a large extent due to this complexity. A company needs to understand the whole system while an attacker only needs find one way in. At the same time, there is a large set of attack types that are utilised and plenty of proposed defence mechanisms. This course main content aims to develop students' understanding of:

- the complex IT landscape of today by creating models of such.
- which attacks that are utilised today to cause harm and how these can propagate through a large network.
- what defences there are and when they are best suited against different attack types.
- how risk can be calculated and used to prioritise security work.

## Intended learning outcomes

After passing the course, the students should be able to:

- model threats in large-scale computer systems (including software, networks etc),
- simulate attacks in large-scale computer systems
- carry out risk analysis based on a model and simulation
- describe which defence mechanisms computer system can have
- report and present models, simulation, risk analysis, and defense strategy for a given system

In order to:

- understand and explain which threats a specific system can have
- understand and explain how attacks work and propagate through a system architecture
- argue why certain risks should be prioritised
- choose the right defence to decrease risk.

## Course syllabus

Today's large complex systems-of-systems are difficult to overlook and manage, while an attacker only needs to find one vulnerability to get in. This course teaches methods for analyzing threats, risks and defense mechanisms of large systems, which can streamline security work and improve protection.

- The schedule can be found online at KTH web. All lectures and seminars will take place online at Zoom incl. the Guest lectures. We will also provide a limited set of seats at Campus.
- Some lectures follow a flipped classroom set-up. Instructions for this can be found below.
- Course material can be found on Canvas incl. recorded lectures and slides and example models.

## Individual project

The purpose of the course assignment is to give skills in, and understanding of, the area of security analysis of large-scale computer systems. More specifically to learn a methodology, threat modeling, to assess cyber security risk of a (large-scale) IT system. The assignment, your project, is designed to be as realistic as possible. This means that not all information needed to solve the assignment is provided in the description. It is therefore necessary for you to make assumptions that are realistic as well as seek new information in order to pass the assignment.

The first thing you need to do is to choose what organisation you will represent and thus do the threat model of. You are free to choose what type of organisation you want, however we encourage you to pick a cyber security incident reported in the media and reverse engineer it. This shall tell you how the organization works and allow you to explore the main attack as well as other attack possibilities.

The course assignments are all carried out individually; this includes:

- the main project,
- the drafts and peer-reviews,
- the guest lectures, and
- the oral presentation.

### Grading

The grading is described below.

The main part is the final report of the project assignment, other graded parts are the oral presentation, attendance at guest lectures, handing in and peer-reviewing drafts in a first iteration.

The five different phases of the method used in the course will be at the core of the examination, that is:

1. business analysis,
2. system definition & decomposition,
3. threat analysis,
4. attack & resilience analysis, and
5. risk assessment & recommendations.

### Submission of assignments

Deliver each assignment in one pdf and name it with your name and assignment tag (e.g. RobertLagerstrom\_finalreport.pdf). Your name must be on the first page of the reports, both draft versions and final version.

For the peer-reviews have your name (the peer-reviewer), the phase, and name of the reviewed (e.g. RLagerstrom\_phase3\_MEkstedt.pdf). Make sure your name, phase, and who was reviewed is presented on the front page.

Only .pdf is accepted as file format.

**Deadlines for handing in the assignments are found on the respective assignment submission blocks in Canvas.**

### References and plagiarism

The main principle behind plagiarism is that you should be responsible for what you submit as your work. Letting others (in particular the teachers) think that something you submitted was your work when in fact it was not, is plagiarism. **It is your responsibility to make sure that no one makes that mistake.** So be clear with references to others when you are using work and ideas from others.

The use of references is mandatory. When you use a fact from some source you should include a reference to this source. Use references according to this or some similar standard but be consistent. For instance:

"Early assessment of system characteristics in software projects is one of the main concerns of the discipline of software architecture [1]."

List of references:

[1] Heineman, G., W. Councill (Eds), *Component-based software engineering: Putting the pieces together*, Addison-Wesley, 2001.

[2] Wikipedia: Enterprise

Architecture, [http://en.wikipedia.org/wiki/Enterprise architecture](http://en.wikipedia.org/wiki/Enterprise_architecture) (Links to an external site.), accessed 2012-03-18

Please note that when solving the project assignments co-operation between individuals is allowed and even encouraged. However, you are responsible for the content of your own reports and **any plagiarism will result in an immediate failing of the assignment in addition to a written report to KTH's central disciplinary committee**. This means that all students should write their own reports. You are not allowed to copy text from other person and you are not allowed to copy text from the Internet. If you want to use a quote from a source, it must be clearly indicated that it is a quote. The reports will be checked with respect to plagiarism using automated scanners.

*Attendance at guest lectures*

Attendance at guest lectures is mandatory, and will be registered on one or more lists that are passed around during the lectures. The attendance requirement is part of the grading (Pass/Fail). Cheating with the attendance list is thus equivalent to cheating on an exam. Attendance is for the whole lecture. **You are not allowed to be late and at the same time sign-up for attendance.**

***For questions and more information about plagiarism and how to avoid it see link on Canvas, post a question in the Discussion section on Canvas or contact the teachers directly.***

Administration

The course teachers are located at Teknikringen 33. The easiest way to contact us is by Canvas or email.

Any complaints regarding the grading of the assignments should be sent to the course teachers no later than one week after the result has been posted.

Disability

If you have a disability, you may receive support from Funka, KTH's coordinator for students with disabilities, see <https://www.kth.se/en/student/studentliv/funktionsnedsattningLinks> to an external site. . Please inform the course coordinator if you have special needs and show your certificate from Funka.

- Support measures under code R (i.e. adjustments related to space, time, and physical circumstances) are generally granted by the examiner.

- Support measures under code P (pedagogical measures) may be granted or rejected by the examiner, after you have applied for this in accordance with KTH rules. Normally, support measures under code P will be granted.

## Flipped classroom

Flipped Classroom describes a concept where the lecture time is not used to present the content that the students should be aware of, but to discuss the content and solve open questions.

The basic concept follows the scheme that the students are provided with material that they have to go through until a certain date. At this date, there will be a lecture where the students can ask questions related to the material. To guarantee that the questions can be answered, we ask the students to provide the questions by 17:00 the weekday day before the lecture. For those questions, we guarantee that they will be answered during the lecture. Other questions can be asked during the lecture as well, but it might be that we will answer them afterwards. If there are no further questions, we are going to end the lecture even if there is still time left.

The flipped classroom concept will be used for all activities listed in the schedule as Q&A sessions.

An exemplary time schedule:

- Day X - N We release a list of material and videos that should be studied by the students.
- Day X – 1 17:00 The students have watched the videos and studied the material. On Canvas, the students write a list of questions that they want to have answered during the lecture the next day.
- Day X The lecture starts (Zoom) and the in-time on Canvas listed questions are discussed. Further questions are answered if possible. The lecture ends when all questions are answered and no additional are brought up.

## Main assignment description - project

### Preliminaries

The purpose of the course assignment is to give skills in, and understanding of, the area of security analysis of large-scale computer systems. More specifically to learn a methodology, threat modeling, to assess cyber security risk of a (large-scale) IT system. The assignment, your project, is designed to be as realistic as possible. This means that not all information needed to solve the assignment is provided in the description. It is therefore necessary for you to make assumptions that are realistic as well as seek new information in order to pass the assignment. (*Realistic* here does not always mean *real*)

### Introduction

You have just been hired as the chief security architect at an enterprise (see below on suggestions). The enterprise's Chief Information Security Officer (CISO), who is also quite new in office, is giving you the assignment to do the annual risk analysis. Today the enterprise has a large number of information systems that provides services to various parts of the business. During the recent years, the systems have been integrated with each other using different integration mechanisms to support various processes, such as sales, marketing, accounting and IT-support.

Unfortunately, the company has lost control over the complete picture of this system-of-systems since it has been, and still is, under constant change. In fact, the company has never really had the complete picture. Every year, new systems are developed and introduced, old systems are extended, modified, integrated with each other, and retired. These changes are

the result of many different stakeholders' requirements and many developers' actions and not of a grand master plan.

The CISO, however, has realized that it is difficult to do a good (quantitative) risk analysis of the company IT infrastructure without knowing what systems they have, how these depend on each other, what data that is flowing, what roles that have access to different parts, what network technologies that are being used et cetera.

The Chief Executive Officer (CEO) and the board of directors have during the last years experienced that it has been hard to make good decisions based on the qualitative and ad-hoc risk analyses they get. Also, the pressure from new laws, increasing digitization, and an increasing number of malicious attack attempts have made them prioritize these questions. And this is where you as the chief security architect come in. You are assigned to do a more thorough risk analysis, one that is quantitative and data-driven, that reflects the business, the IT environment, and the current threats, so that the CISO and people responsible for making strategic decisions have an up-to-date understanding of the current situation. The risk analysis should be created in order to support the CISO.

## **The Main Assignment**

### **Individual work**

#### **Mandatory**

*(For the assignment, the first thing you need to do is decide what enterprise you have been hired at. We encourage you to pick a known cyber attack incident and read up on it as has been publicly reported and from this material deduct/reverse engineer how the organization works and make additional assumptions about the organization. Then you would also have a very specific attack to study as a starting point. e.g., the [SolarWinds \(Links to an external site.\)](#) hack in 2021, or what happened at [Capital One \(Links to an external site.\)](#) in 2019, or at [Yahoo \(Links to an external site.\)](#) in 2013, or the [Ukraine \(Links to an external site.\)](#) power grid in 2015? Other sources of inspiration include incidents reported in the report by [MSB \(Links to an external site.\)](#). Another option could be to choose an organization you are already comfortable with, maybe you have work experience you can have use of, or an organization that you want to learn more about and thus have a motivation to dig deeper into. )*

You know that threat modeling is appropriate to assess the security level of the enterprise and its applications. Further, it allows identifying weaknesses within the architecture. With this approach you can visualize and concretize what the risks are and how these could be mitigated.

Fortunately, you have heard about a method called Yacraf that that seems to fit your needs perfectly.

As the **main delivery**, the CISO is expecting a report consisting of five main parts;

- 0) scope & delimitations,
- 1) business analysis,
- 2) system definition & decomposition,
- 3) threat analysis,
- 4) attack & resilience analysis, and
- 5) risk assessment & recommendations.

You are free to use any tools you like to support the work in the different phases of the assignment.

Using graphical models, tables, and such can solve many steps. However, these can't stand by themselves, but need to be complemented with text explaining the figures and your

analysis. Furthermore, you are creating a report, and as such, some text is likely needed to guide the reader throughout the report, i.e. understandability is important, thus it must neither be too big and complicated nor too small and trivial.

### **Assignment assessment and grading criteria**

The assignment will be evaluated according to the criteria listed under grading below. (*This means that understanding these criteria as soon as possible is essential for succeeding with the assignment. As your work progresses, carefully cross-examine your work with respect to the criteria.*)

Mapping to intended learning outcomes

#### **Phase 1: Business value of system**

- Loss events (breach impact) based on business architecture (use cases) and business goals  
\*\*\*Model threats in large-scale computer systems\*\*\*

#### **Phase 2: System definition and decomposition**

- Data flow diagrams based on system assets, actors, accounts, and authorization  
\*\*\*Model threats in large-scale computer systems\*\*\*

#### **Phase 3: Threat analysis**

- Abuse cases based on attacker profiles  
\*\*\*Simulate attacks in large-scale computer systems\*\*\*

#### **Phase 4: Attack and resilience analysis**

- Attack trees based on vulnerabilities  
\*\*\*Simulate attacks in large-scale computer systems / Describe which defense mechanisms computer system can have\*\*\*

#### **Phase 5: Risk assessment and recommendations**

\*\*\*Carry out risk analysis based on a model and simulation / Describe which defense mechanisms computer system can have\*\*\*

[First iteration drafts for peer-review](#)

### **Individual work**

#### **Mandatory**

#### **Description**

Before delivering the final report the CISO would like to feel confident that you are on the right track and have thus asked you deliver two drafts. This enables the CISO to guide you in the right direction for the final report (on which both your future careers depend). Moreover, to really ensure the quality of the work the CISO has ordered a few of your colleagues in the security group to review it (peer-reviewing).

The main goal with this assignment is to familiarize yourself with the threat modeling method used and what kind of challenges each phases of it include, as well as the company you are modeling. You need to collect some general material or other experiences and conjecture how the chosen type of company typically works.

There will be two iterations with one draft handed in and peer-reviewing for each.

The initial drafts does not have to be large or complex but comprehensive enough to show your general direction, it will mainly consist of models, figures, lists, et cetera and not complete text (but some supporting text where needed). But of course, better drafts usually means more relevant feedback. As you will continue working on the drafts, its quality will not impact your final grade. These submissions are for the purpose of directing your continued work.

Each person will peer-review other peoples drafts for each iteration.

After each peer-review round you need to write a short 4-7 sentence "This is what I learned during the peer-review process"-report and hand in.

The pedagogical idea with having the drafts and peer-reviews early in the course with a rather tight time schedule is twofold: 1) We want you to get a flying start by thinking and reading about all the phases early on. You learn a lot going through a full iteration of the method. 2) We want you to have plenty of time left for the second iteration to complete the project assignment. Thus being able to implement all the the things you have learned from the first drafts and peer-reviews (and all the other activities in the course).

### **Evaluation Criteria**

The drafts and peer-reviews are **mandatory** and graded **Pass/Fail**.

Everyone must hand in their drafts on Canvas.

Then either

- you read the two reports assigned to you and provide written feedback on each of the two reports

or

- you attend the peer-review seminar and provide oral feedback live during the seminar. For those attending the peer-review seminar you don't have to read the, in Canvas assigned, reports (these will most likely be different then the ones you review in the seminar). Instead you will get the chance to read drafts during the seminar, listen to others presenting their drafts, and provide feedback during the seminar.

Everyone needs to write the short 4-7 sentence "This is what I learned during the peer-review process"-report and hand in.

### [Final Presentations](#)

#### ***To be performed individually***

#### ***Mandatory***

#### **Description**

One of the key success factors for a long-lived threat modeling initiative is that it is supported by senior management as well as by the operative staff at the business units and the people in the IT organization. One of the most important assignments for chief security architects is to communicate and explain the purpose of threat modeling as well as the ongoing and future work within the area. Your assignment is to hold a 10-minute executive presentation about the current threat and risk status, as well as your proposed mitigations, for the employees of the enterprise (i.e. the course participants) as well as the CISO and the CEO (the teachers). This presentation will serve as the beginning of an era of more structured and efficient cyber security risk management at the enterprise.

It is important that the audience after the presentation has understood the following:

- Background information, explaining the need for the presented work. (This is an important thing to communicate and synchronize within enterprises; a common goal!)
- How the business works, and your business impact analysis.
- What is the IT support offered to the business or, vice versa, how dependent on IT are the various parts of the business? What does the system architecture look like.
- What are the most important threats and attack profiles you worry about at your enterprise?
- What kind vulnerabilities did you find and what type of attacks did that enable?

- What does your risk analysis result in?
- Which mitigations do you suggest based on the risk analysis? Explain your reasoning.

Focus on the particulars of your case and company and not too much on the underlying methodology. The employees of the company are more interested in the results and what to do rather than exactly how you got there (especially the parts that are the same for others doing similar work).

**Note!** The presentation must be prepared so that it can be **given in English**. If only Swedish-speaking people are present at the presentation seminar, Swedish is also OK.

### Evaluation Criteria

The oral presentation is **mandatory**. The presentation is **Pass/Fail** graded. The presentations will be evaluated on the communicative performance, i.e. on the extent to which the contributions are correctly and efficiently communicated to the audience. This will typically be affected by the structure and form of the presentation and slides, the clarity of the argumentation, the presenter's oral performance, the timeliness of the presentation, and on the responses to questions posed by the audience.

### Grading

To pass the course (E) you need to pass the project assignment (6.0 credits) and the seminar series (1.5 credits), this includes:

- Passing the project report (evaluated based on properties listed below), A-F
- Pass the oral presentation, P/F
- Attend the guest lectures, P/F
- Hand in peer-review learning reports, P/F

The project report will be the main source of evaluation and it will be assessed based on a set of properties, namely; consistency, correctness, coverage, variance, (technical) detail, realism & motivation, and balance/assurance. These are described and exemplified in the table below.

Moreover, the report must be well-structured, understandable and readable. E.g. it must be free of spelling errors, inconsistent or incomprehensible sentences, and grammatical errors. It must have reasonable size of paragraphs and sentences. It must use numbered headings and it must be easy to follow all internal references in the report. It must use figure/table text and reference figures/tables in text (e.g. figure text description: "Figure 1. A business model". In text: "In Figure 1 the business model is described" ).

Finally, the **report must be handed in on time**.

Property	Evaluation criterion	Comment and examples
Consistency	The models follow a single and structured underlying framework(metamodel), preferably the one presented in the course. And the terminology is used correct and consistently.	E.g. are multiplicities followed and are phenomena modeled according to the logic of the language? E.g. "threat," "risk," and "vulnerability" should mean different things (unless something else is explicitly stated).



	<p>The models are consistent within each phase (sub model) and between one another (across the phases).</p>	<p>E.g. are the same assets modelled in the various DFDs? Are the vulnerabilities discussed relate to the assets in the system architecture, the attack vectors/trees relate to the identified vulnerabilities and the assets, etc. It must be possible to follow/track/relate all model elements to one another. Everything needs to be consistent. The consistency should be demonstrated (or rather the report should not leave places where the consistency can be doubted or questioned).</p>
Correctness	<p>The risk calculations, conclusions, and other syntheses are made correctly.</p>	<p>The fusion of parameters should follow the described method or otherwise make sense/be motivated. E.g. it doesn't make sense to sum the attacker capability with a vulnerability score. And the calculations must match the developed models.</p>
Coverage	<p>The models are covering some well-defined problem domain in a reasonably complete way.</p>	<p>E.g. for attack models - apparent attack vectors should not be missing, for system models - key components should not be missing, for business goals - obvious losses should not be missing. Unless these parts are explicitly stated as delimitations, then it could be ok with less coverage.</p> <p>The chosen framework is fully covered.</p>

Variance	The models cover a variety of phenomena.	<p>Is the chosen resolution appropriate?</p> <p>Not all models/analyses should follow the same underlying logic. Say three web servers with very similar vulnerabilities. One should demonstrate ability to analyze a wide variety of security topics. Different types of web vulnerabilities, but even better also network vulnerabilities and others.</p>
(Technical) detail	<p>The models should be on such a depth so that they are not trivial or superficial. Also, it should be clear what the technical challenges /implementations / weaknesses /attacks are.</p>	<p>E.g. a non-technical attack vector is DoS web server &lt;- deploy botnet &lt;- buy botnet. Remedy: buy load balancer.</p>
Realism and motivation	<p>The model input data is realistic. Some things are straight forward and not questionable while other assumptions and interpretations motivation and external references are needed to back up the claims of the models.</p>	<p>In principle there are two ways of motivating non-questionable claims, either (trustworthy) references are put forward as support or you put forward your own argumentation that is convincing the reader about the realism. This can relate to anything from costs of loss to common system architecture solutions and technology stacks, to threat profiles and attack vectors.</p> <p>Strong references, and use of, external sources. E.g. we have analyzed some certain MITRE ATT&amp;CK techniques, or some vulnerability class. E.g. the use of threat emulation profiles.</p> <p>For the threat analysis there are many parameters that is qualitatively assessed (e.g. prob. of action). For these</p>

Balance/assurance

The confidence of the recommendation analysis and conclusions are declared and discussed.

assessments some form of reasoning / motivation is needed.

Motivating why you are not diving deeper into something is a good thing. E.g. we consider TLS 1.3 secure, because ...

E.g. this relates to making probabilistic risk calculations, but also the intermediary focus of the threat modeling per se. E.g. it is decided to focus on a certain abuse case or vulnerability. How sure are we that this focus/delimitation is the correct one? Balance the level of detail so that not one part/aspect of the models get a lot of attention and others little unless this difference in importance is clearly motivated. This criteria also relates to the remaining uncertainty after the realism has been motivated.

All seven properties will be assessed and all must pass a minimum bar. If one criterion is assessed and considered as failed the grade Fx will be used in order for you to re-do that part to pass the course.

Each of the eight criteria will be evaluated using a scale of Fail (-1), Sufficient/Pass (0), Good (1), Very good (2), and Excellent (3). Thus, a maximum of 24 points can be gathered. Grades will be set based on the following schema:

A = 20-24

B = 15-19

C = 10-14

D = 5-9

E = 0-4