AWS re:Invent

WPS321-R

Managing InfoSec risk during cloud adoption

Stephen Alexander
Senior Solutions Architect
Amazon Web Services

AWS re:Invent

aws

Agenda

The state of enterprise risk assessment

A different approach: Risk analysis using latency curves

Risk analysis using latency curves: A walk-through

Conduct your own risk analysis

The state of enterprise risk assessment

"The winning general is the one who can best act on imperfect information and half-formed theories."

Napoleon Bonaparte
Emperor of France

Customers often ask

How risky is it to move a specific workload to AWS?

Is our data as safe on AWS as it is in our on-premises data center?

Can we maintain the same uptime on AWS as in our data center?

How does the security in our data center compare to security with AWS?

How can we be certain our security controls are working on AWS?

Risk management process

1. **Identify** risks reviewing authoritative risk sources

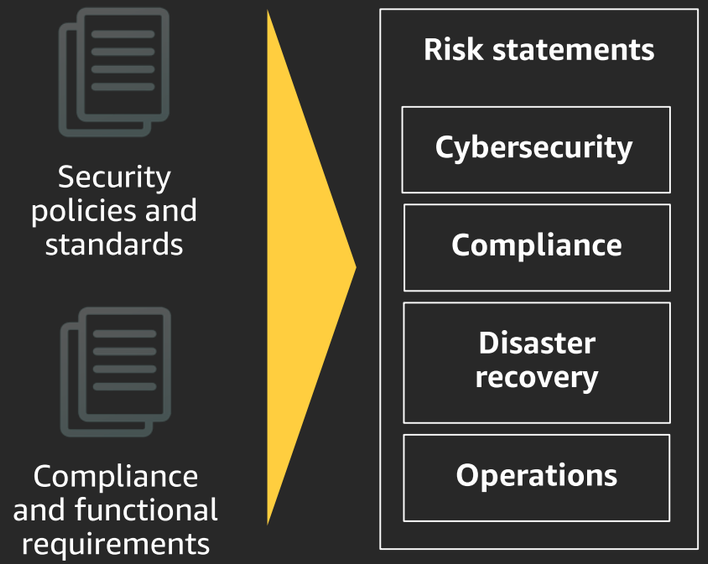

2. **Analyze** risks for more informed decision-making

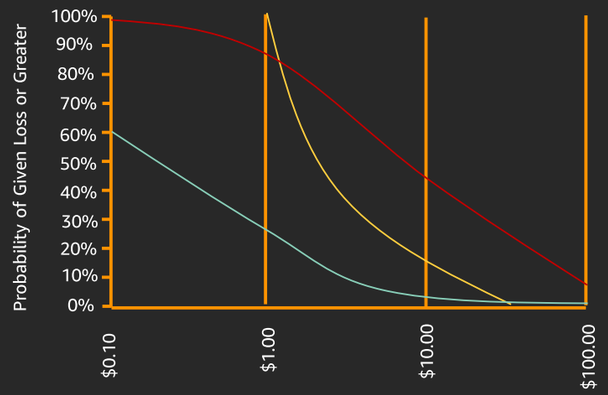

3. **Develop** risk responses and assign ownership

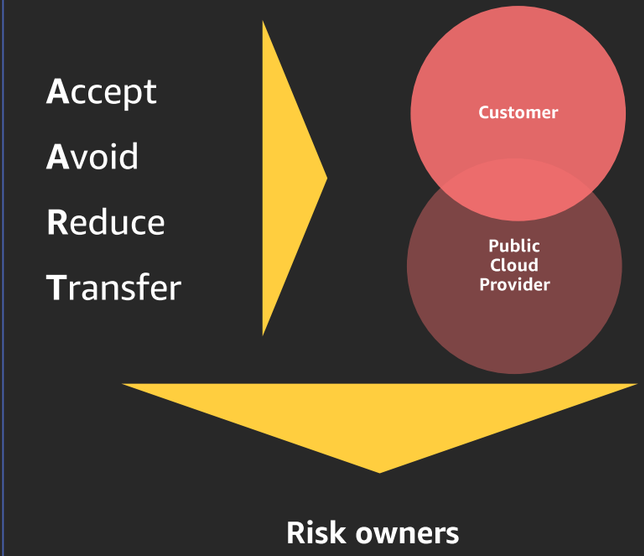

Key outputs

- Immediate mitigations
- Compensating controls
- Mitigation roadmap

Three typical risk analysis (step 2) outputs

Actuarial table

Table 2000CM

Age x	l_x	Age x	l_x	Age x	l_x
0	100000	37	96921	74	66882
1	99305	38	96767	75	64561
2	99255	39	96600	76	62091
3	99222	40	96419	77	59476
4	99197	41	96223	78	56721
5	99176	42	96010	79	53833
6	99158	43	95782	80	50819
7	99140	44	95535	81	47694
8	99124	45	95268	82	44475
9	99110	46	94981	83	41181
10	99097	47	94670	84	37837
11	99085	48	94335	85	34471
12	99073	49	93975	86	31114
13	99057	50	93591	87	27799
14	99033	51	93180	88	24564
15	98998	52	92741	89	21443
16	98950	53	92270	90	18472
17	98891	54	91762	91	15685
18	98822	55	91211	92	13111
19	98745	56	90607	93	10773
20	98664	57	89947	94	8690
21	98577	58	89225	95	6871
22	98485	59	88441	96	5315
23	98390	60	87595	97	4016
24	98295	61	86681	98	2959
25	98202	62	85691	99	2122
26	98111	63	84620	100	1477
27	98022	64	83465	101	997
28	97934	65	82224	102	650
29	97844	66	80916	103	410
30	97750	67	79530	104	248
31	97652	68	78054	105	144
32	97549	69	76478	106	81
33	97441	70	74794	107	43
34	97324	71	73001	108	22
35	97199	72	71092	109	11
36	97065	73	69056	110	0

Heat maps

Impact	1	2	3	4	5
Probability	Negligible	Minor	Moderate	Significant	Severe
81-100%	Low Risk	Moderate Risk	High Risk	Extreme Risk	Extreme Risk
61-80%	Minimum Risk	Low Risk	Moderate Risk	High Risk	Extreme Risk
41-60%	Minimum Risk	Low Risk	Moderate Risk	High Risk	High Risk
21-40%	Minimum Risk	Low Risk	Low Risk	Moderate Risk	High Risk
1-20%	Minimum Risk	Minimum Risk	Low Risk	Moderate Risk	High Risk

Latency curves

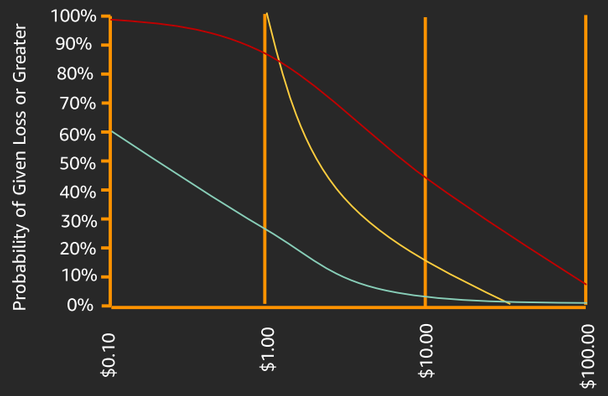

Given loss or greater (millions)

How many organizations analyze InfoSec risk

Typical risk register

Risk	Description
1	Disclosure of proprietary customer data
2	Disclosure of quarterly financials in quiet period
3	Disclosure of CEO email due to spear-phishing
4	Catastrophic power failure takes data center offline for > four (4) hours
…	
N	Godzilla crushes data center

Typical risk heat map

Impact	1	2	3	4	5
Probability	Negligible	Minor	Moderate	Significant	Severe
81-100%	Low Risk	Moderate Risk	High Risk	Extreme Risk	Extreme Risk
61-80%	Minimum Risk	Low Risk	Moderate Risk	High Risk	Extreme Risk
41-60%	Minimum Risk	Low Risk	Moderate Risk	High Risk	High Risk
21-40%	Minimum Risk	Low Risk	Low Risk	Moderate Risk	High Risk
1-20%	Minimum Risk	Minimum Risk	Low Risk	Moderate Risk	High Risk

Heat maps disguise bias

2005: A survey of NATO officers demonstrates that “highly likely” means anywhere between 40% and 100% likely (Heuer)

2006: Studies find experts choose “1” more often on a scale of “1-10” regardless of the subject matter the number supposedly represents (Rottenstreich)

2009: A survey of students and faculty demonstrates that “very likely” means anywhere between 43% and 99% likely (Budescu)

2016: When percentages are explicitly defined (e.g., “highly likely” represents 90-99%), survey participants violate the rules 52% of the time (Hubbard/Seiersen)

2016: A survey of experts using the ordinal scale “1-5” finds that they choose the values “3” or “4” three times as often as the other values (Hubbard/Seiersen)

Heat maps also do not convey any information as to when an event may occur (next month, this year, next two years, next decade)

A different approach: Risk analysis using latency curves

Let’s start with some definitions

- **Risk:** The probable frequency and probable magnitude of future loss
- **Risk model:** The calculation of risk that a specific threat and a given threat scenario pose to an asset
- **Threat actor:** A causal agent (person, entity, or natural source) that, together with a specific threat scenario, is the source of risk to an asset; threat actors can be adversarial or accidental
- **Threat scenario:** A series or combination of events that, together with a threat actor, poses risk to a specific asset; a scenario represents the dynamic execution of a *tactic*
- **Asset:** The target of the threat scenario, such as data, a material object, or a person

"Don't think of probability or uncertainty as the lack of knowledge. Think of them instead as a very detailed description of what you do know."

Ronald Howard
Professor in the Department of Engineering-Economic Systems
Stanford University

A new and improved risk register

Redefined risk register

Risk	Description
1	What is the probability of **material** disclosure of proprietary customer data in the next two years?
2	What is the probability of **material** disclosure of quarterly financials in a quiet period anytime over the next four quarters?
3	What is the probability that disclosure of the CEO's email due to spear-phishing will have a **material** effect on the company?
4	What is the probability that a catastrophic power failure will take the data center offline for > four (4) hours anytime in the next three years (before we move 100% to the public cloud)? (This is 100% material.)
…	
N	What is the probability that Godzilla will crush the data center in the next three years (before we move 100% to the public cloud)? (This is 100% material.)

Typical risk heat map

Impact	1	2	3	4	5
Probability	Negligible	Minor	Moderate	Significant	Severe
81-100%	Low Risk	Moderate Risk	High Risk	Extreme Risk	Extreme Risk
61-80%	Minimum Risk	Low Risk	Moderate Risk	High Risk	Extreme Risk
41-60%	Minimum Risk	Low Risk	Moderate Risk	High Risk	High Risk
21-40%	Minimum Risk	Low Risk	Low Risk	Moderate Risk	High Risk
1-20%	Minimum Risk	Minimum Risk	Low Risk	Moderate Risk	High Risk

Given what I know, I estimate a 0.0000000001% probability that Godzilla will crush the data center in the next three years (unless the data center is in Tokyo, in which case I estimate a a 0.0000000002% probability)

An example latency curve

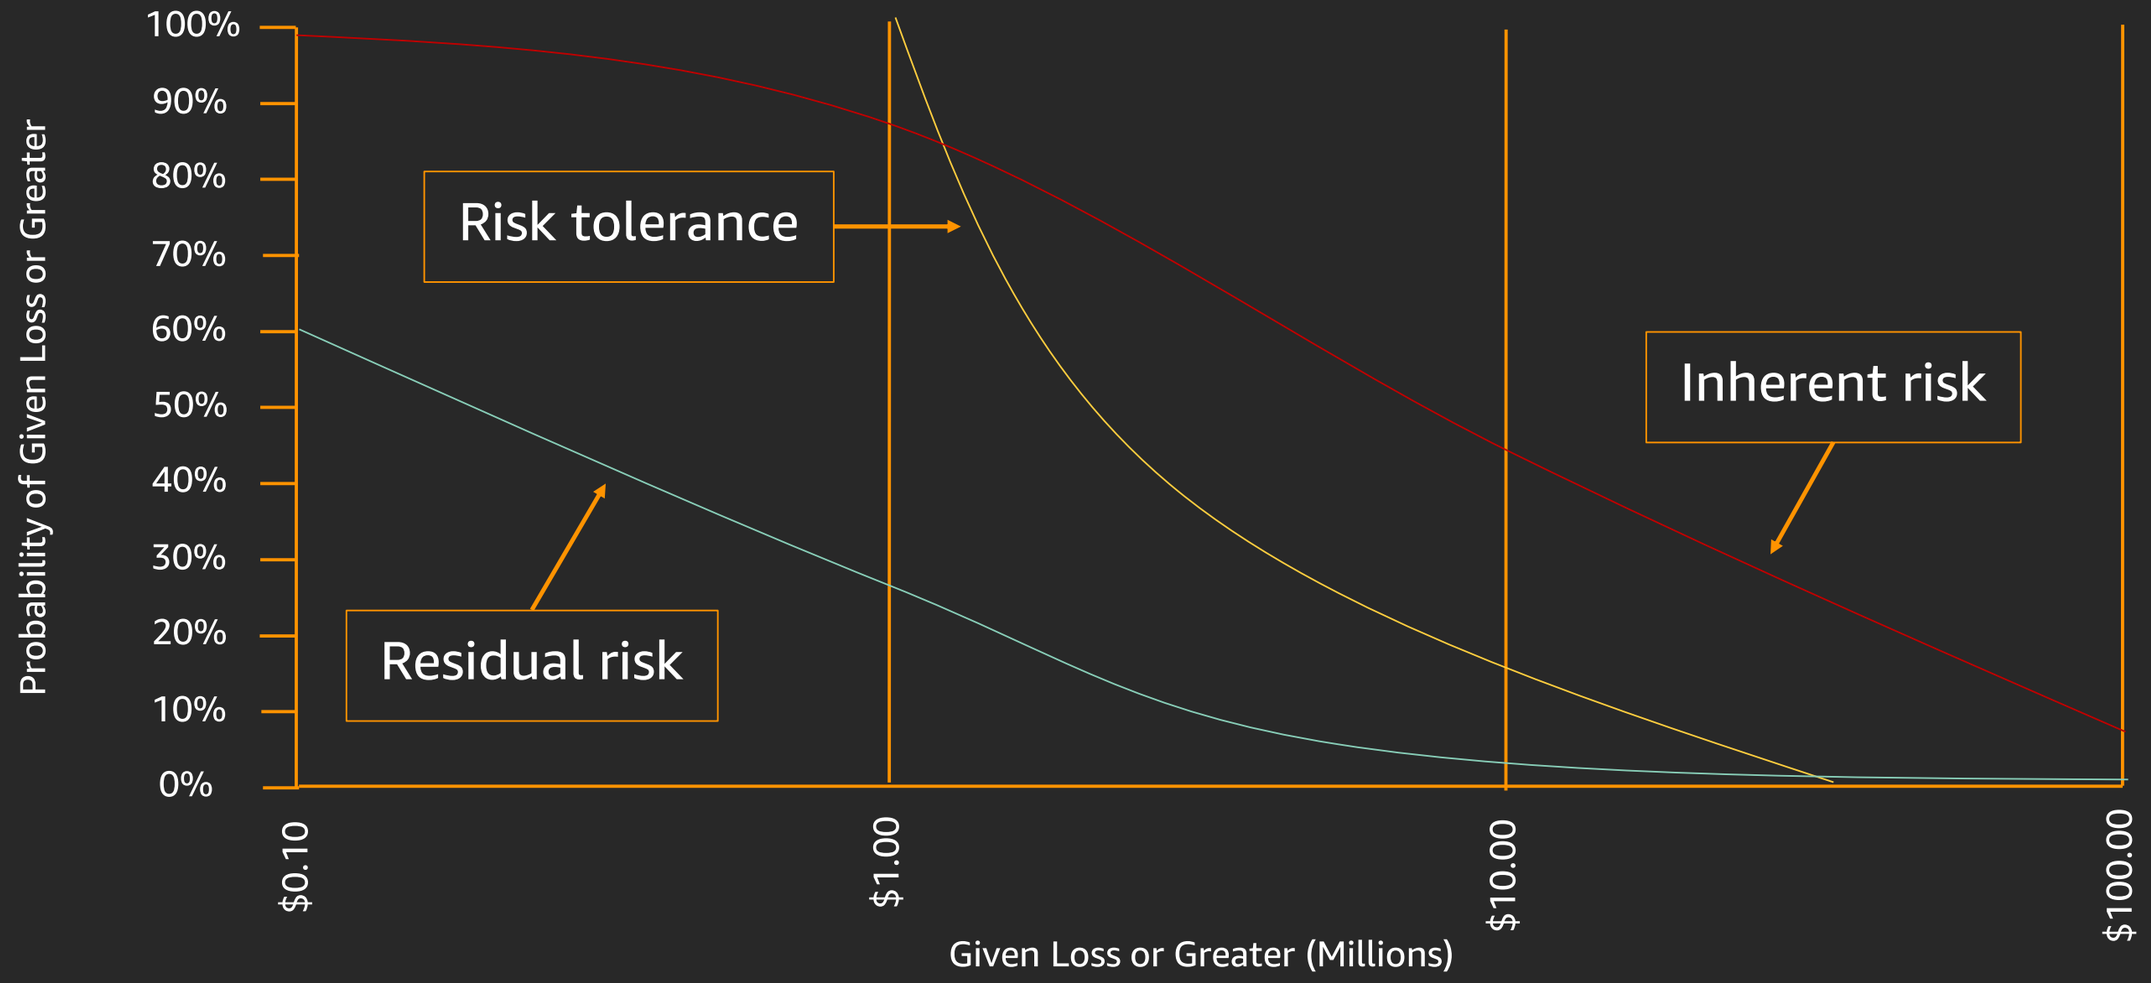

Risk analysis using latency curves: A walk-through

AWS re:Invent

aws

Step 1: Create AWS-specific scenarios

Pre-AWS risk register

Risk	Description
1	What is the probability of **material** disclosure of proprietary customer data in the next two years (***if that data remains where it is currently stored***)?
2	What is the probability of **material** disclosure of quarterly financials in a quiet period anytime over the next four quarters (***assuming that these financials remain on premises***)?
3	What is the probability that disclosure of the CEO's email due to spear-phishing will have a **material** effect on the company?
4	What is the probability that a catastrophic power failure will take the ***data center offline*** for > four (4) hours anytime in the next three years (before we move 100% to the public cloud)? (***This is 100% material if no DC redundancy.***)
…	
N	What is the probability that Godzilla will ***crush the data center*** in the next three years (before we move 100% to the public cloud)? (***This is 100% material if no DC redundancy.***)

AWS migration risk register

Risk	Description
1	What is the probability of **material** disclosure of proprietary customer data in the next two years (***if we move that data to AWS***)?
2	What is the probability of **material** disclosure of quarterly financials in a quiet period anytime over the next four quarters (***assuming that these financials move to AWS***)?
3	What is the probability that disclosure of the CEO's email due to spear-phishing will have a **material** effect on the company?
4	What is the probability that a catastrophic power failure will ***take AWS Availability Zone offline*** for > four (4) hours anytime in the next three years? (***This is 100% material without AZ redundancy.***)
…	
N	What is the probability that Godzilla will ***crush AWS AZs*** in the next three years? (***This is 100% material without AZ redundancy.***)

Step 2: How do we estimate?

How do we decide on a simple loss distribution?

Base rate

- Use customer metrics (e.g., data center historic downtime)
- Use industry information (e.g., Ponemon “Cost of a Data Breach,” Gartner, Forrester, MITRE data)

We can combine different individuals’ knowledge, creating different Monte Carlo simulations and then aggregating

Fermi estimation

- “How many piano tuners are in Chicago?”
- Fermi estimation is a process to determine information as accurately as possible within the boundaries of uncertain knowledge
- Think of it as a formal process for educated guesswork

Step 3: Create a simple loss distribution

Loss amount ($M)	Material loss in given scenario
0	–
>0 and <=1	0.6
>1 and <=2	0.2
>2 and <=5	0.1
>5 and <=10	0.1

Step 4: Run Monte Carlo simulation

Total loss	Sample count	Sample percentage
0	11	0.00071
1	75	0.00485
2	181	0.011704
3	342	0.022114
4	489	0.03612
5	546	0.035306
6	630	0.040737
7	749	0.048432
…		

Risk response

Accept

Avoid

Reduce

Transfer

Step 5: Decide on controls

- **Inherent risk:** Given a specific risk register item, the risk that is based on current controls if the scenario occurs
- **Risk tolerance:** Loss the customer is willing to accept based on the probability that the scenario occurs
- **Controls:** The high-level controls that AWS recommends to mitigate inherent risk (e.g., usage of multiple AZs or multiple Regions to protect against downtime)
- **Cost of controls:** How much will the recommended controls cost?
- **Residual (AWS) risk:** If I implement the controls AWS recommends, what does the risk profile look like?

The resulting latency curve

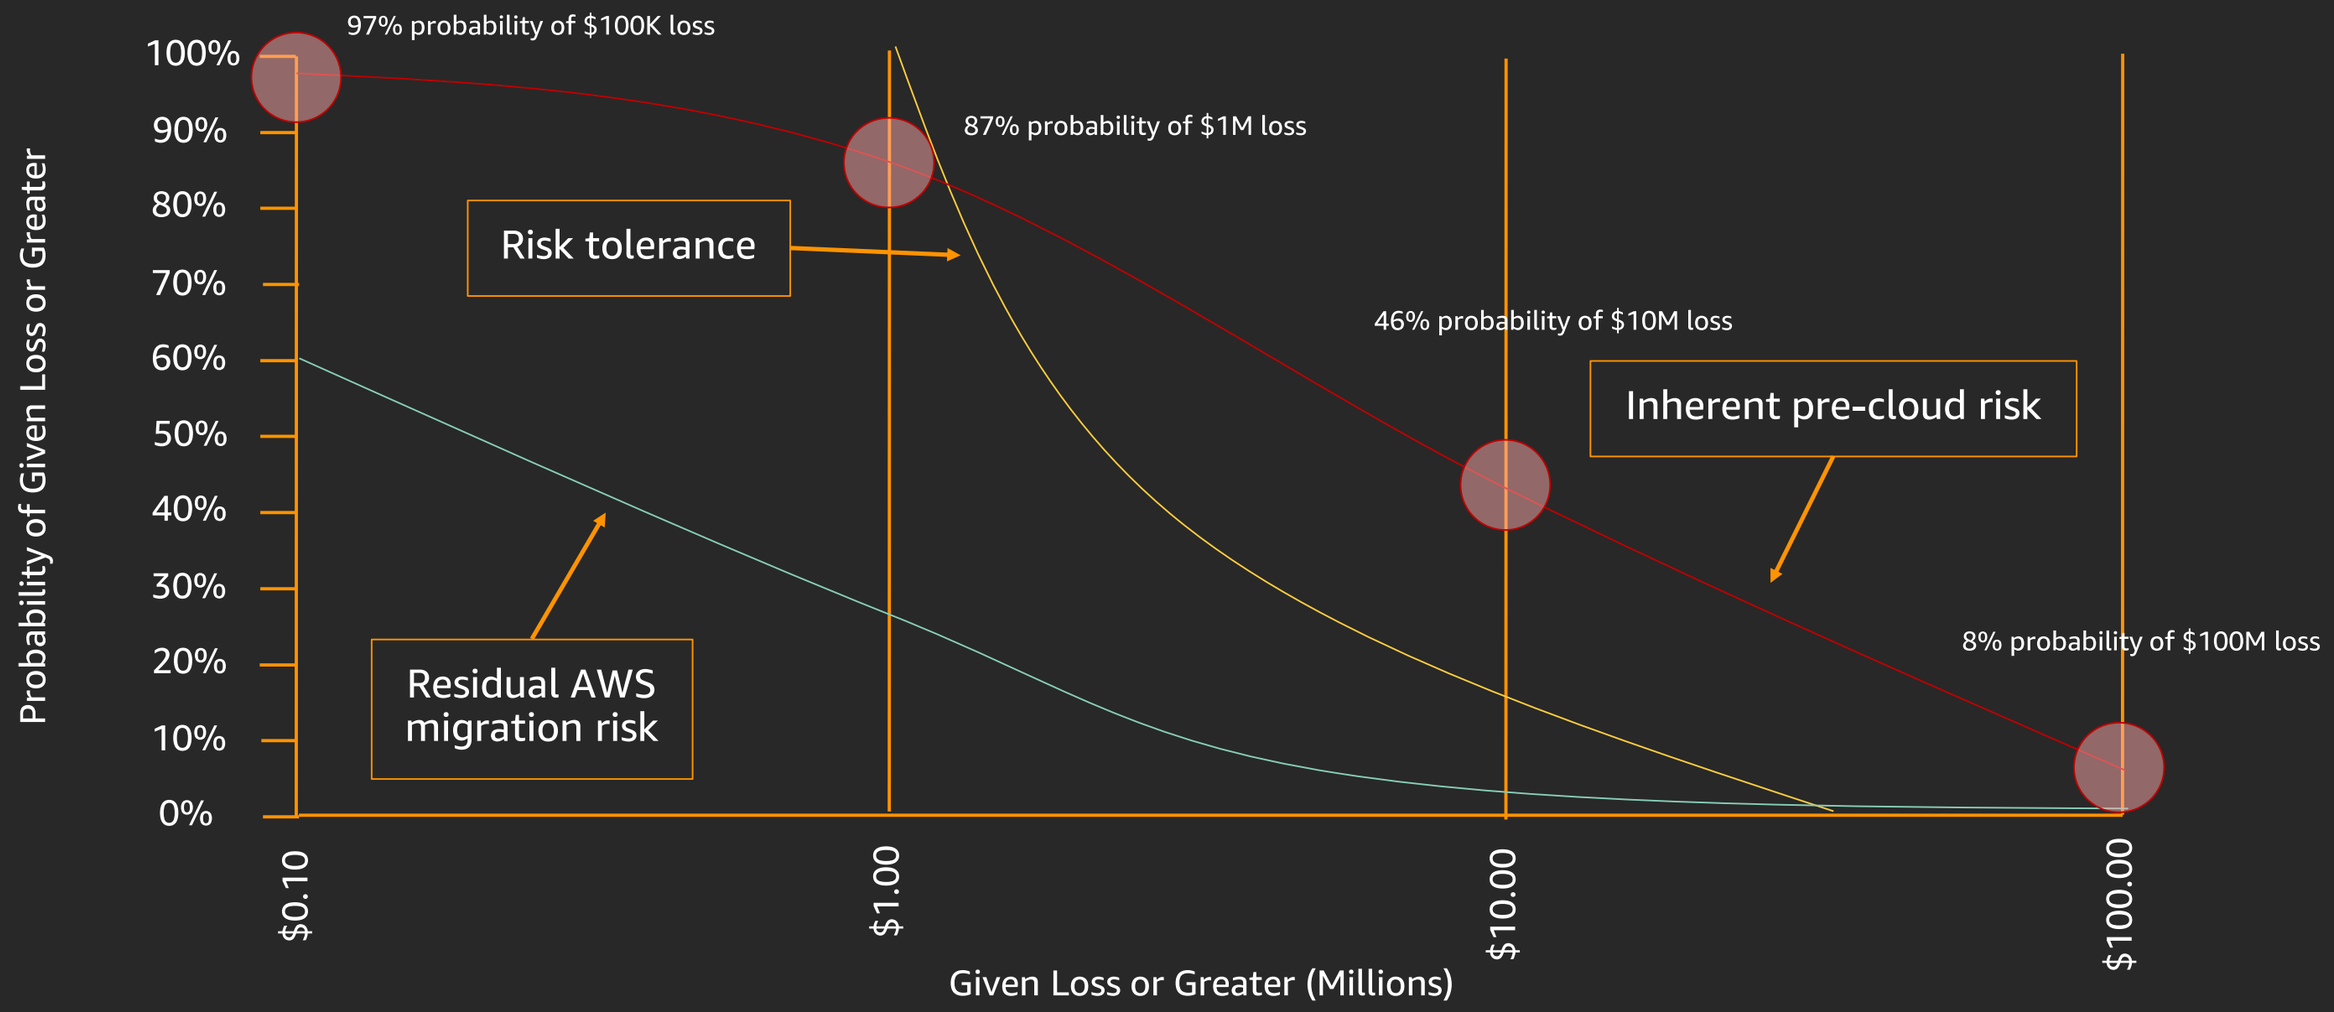

Conduct your own risk analysis

Scenario

You are a key risk manager at Example LLC, an international private equity firm based in New York. Your yearly revenue was $1.2B last fiscal year. Your main US data center is located in northern New Jersey, and you replicate some key servers (80 of your 600 servers) to another data center near Philadelphia. Most importantly, you replicate your key financial data to the secondary data center. Executive leadership and the board want to move to AWS for cost savings, but they want to understand the risks associated with the migration. Their key concerns:

- Is our financial data *as safe* on AWS as it is in our data center?
- Can we maintain the *same uptime* on AWS as in our data center?

Background information and assumptions

- Your IT security budget has been under pressure for the past five years—you spend approximately 20% less, adjusted for inflation, than you did five years ago; you were reading a report that shows that your budget is about the median for organizations your size
- You were massively hit by WannaCry two years ago, affecting both your Windows servers (70% of your servers) and end-user workstations; since then you've been more diligent about patching, but your vulnerability scans still show that your *average time to patch* lags at least three days, even for critical vulnerabilities

Background information and assumptions

- You know of at least three instances over the past year of insiders leaving the company and taking confidential information with them. You have server and workstation DLP, but it is nearing end of support. You do not have a 24/7 SOC—IT security staff receive critical alerts from your SIEM and investigate as needed, but there is no defined SLA.
- Your key financial data is stored encrypted in an Oracle database. You do not have an on-premises HSM.
- In addition to the disruption from WannaCry two years ago, you suffered a cascading outage in your main data center six months ago due to bad weather. The outage was brief (two hours), but recovery took most of the workday.

Analyze your inherent and residual risk!

1. Work together to rewrite your risks to include materiality and temporality

2. Collectively use the provided Excel spreadsheet to estimate a simple loss distribution for the current situation

3. Decide your organization's risk tolerance

4. Collectively use the provided Excel spreadsheet to estimate a simple loss distribution for the AWS recommended solution outlined in the handout

Thank you!

AWS re:Invent

aws

Please complete the session survey in the mobile app.