How Much Is Enough?

How much engineering? How much security? For which kinds of systems? And why?

Andrew Ginter
VP Industrial Security, Waterfall Security

2026 Threat Report – Ransomware

57 Breaches – in 2025, from Jaguar / LandRover to medical nuclear isotopes

"Abundance of caution" – not confident of strength of OT security program

Dependencies – shut down because physical ops need crippled IT systems

Direct OT attack – shut down ops because automation is compromised

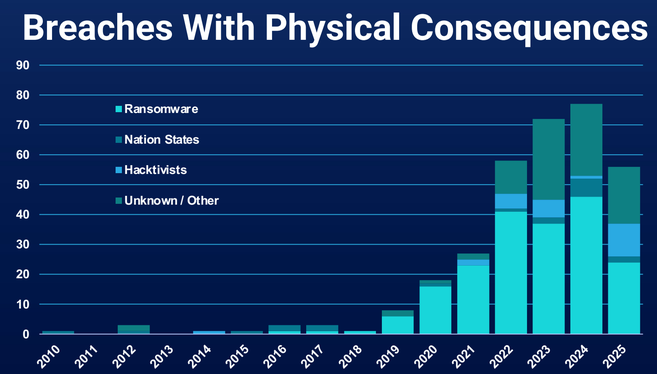

Supply chain – supplier or customer is compromised, and so victim shuts down

Kuala Lumpur Airport – flight delays 2 days because of baggage handling

Attacks overall are down 26% because ransomware attacks are down

Hacktivists and Nation States

Breaches – doubled over previous year

Nation-state-sponsored hacktivists – boundary is blurring – especially in kinetic conflicts

Iranian air defenses – crippled by US military during assault on nuclear program

Russian foodstuffs – cloud tools for tracking animal products crippled – food processing slowed

MSC Antonia – ran aground – GPS spoofing

Nation states are cheap – with more powerful tools in their pockets than they have shown us to date – include a large margin for error for credible nation-state targets

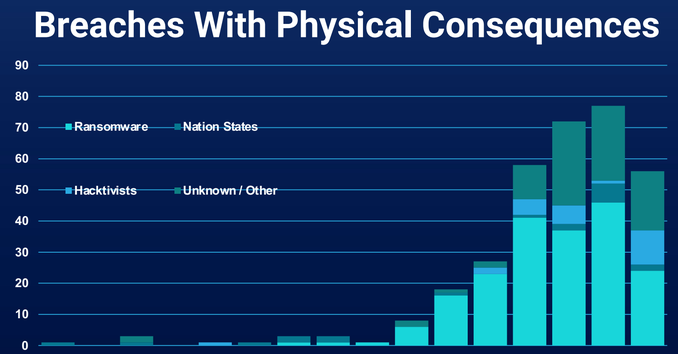

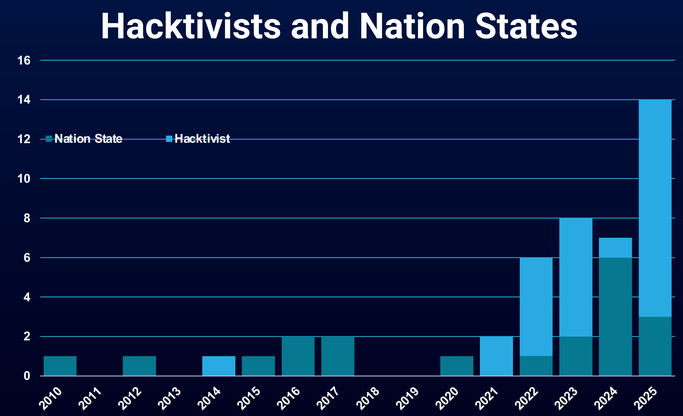

Zero Days Are Upon Us

Anthropic's Claude Mythos – has proven very effective at discovering zero days – eg: Firefox issued an update for 320 previously unknown vulnerabilities

What's different this time? – we faced a similar problem a decade ago with the invention of "fuzzing" tools finding many zero days.

Mythos chains zero-days – taking many apparently low-severity vulnerabilities and combining them into powerful exploits

AIs are automating attacks – increasing attacker productivity 5-10x

Fully-autonomous attackers will have the advantage on OT defenders

Closely track the development of autonomous attack capabilities and invest in deterministic defenses

Attacks Pivot – Use Compromised Machines to Attack

Compromised Machines – "phone home" and are operated manually by remote control

Use Compromised Machines – to attack other machines – in the same network and through firewalls

Breaking Firewalls –

- Shoulder-surf admin password – add "allow all" rule
- Exploit known firewall vulnerability – before it is patched
- Find a zero-day – no patch avail
- Phish a passwd – to any service exposed through the firewall
- Exploit OT application vuln – before it is patched
- Exploit permissions – take over IT Active Directory & create own accounts – log in thru firewall

After all – ransomware criminals Pivot routinely through IT-Internet firewalls to reach IT, don't they?

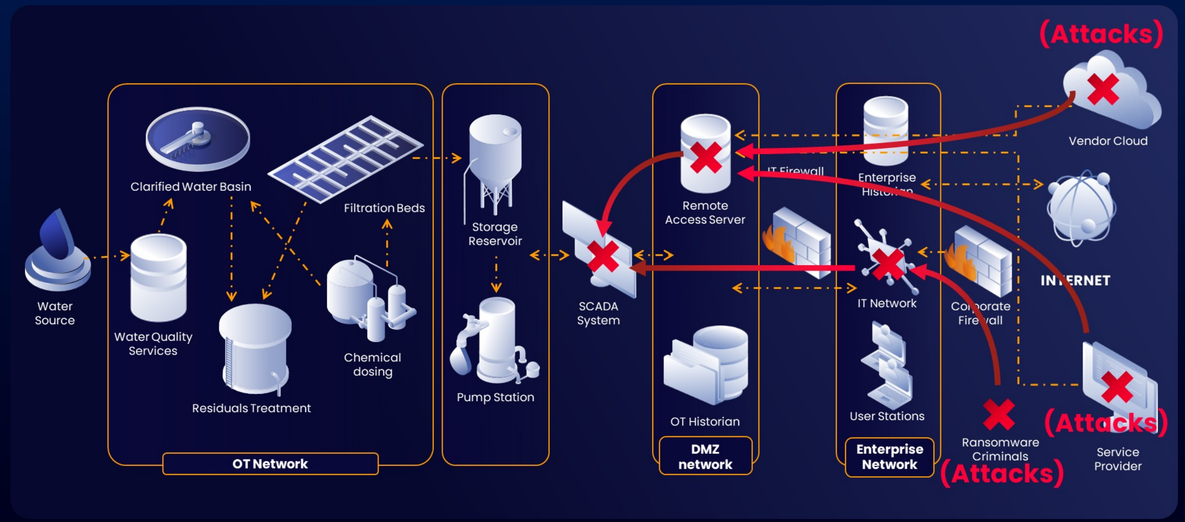

Brief History – OT Security Is Evolving

First gen – like IT – information is the asset – protect the CIA / AIC of the information – "encrypt everything!"

Second gen – physical operations / assets – protect safe, reliable and efficient physical operations. Information is the threat – protect physical ops *from* (attack) information

Engineering Change Control (ECC) – every change risks errors and omissions - study changes first – the more consequential the system, the more engineering we do

Resilience – design risk out of our systems, ride through attacks, maybe at reduced capacity or efficiency, reduce downtime to acceptable levels – non-cyber safeties + detect / respond / recover for reduced downtime

Resilience until recently was more "headline" than "recipe"

How Much Is Enough?

What is risk? Risk is the effect of uncertainty on mission objectives
Our obligation – reasonable responses to credible threats
Universal approximation – ***risk = impact x likelihood***

Likelihood						
Very high	Low	Low	Moderate	High	Very High
High	Low	Low	Moderate	High	Very High
Moderate	Very Low	Low	Moderate	Moderate	High
Low	Very Low	Low	Low	Low	Moderate
Very Low	Very Low	Very Low	Low	Low	Low
	Very Low	Low	Moderate	High	Very High

Level of Impact

Qualitative Impact / Likelihood Rankings

Risk = impact x likelihood – but does 1x5 = 5x1

Likelihood						
	Very high	**Low**	Low	Moderate	High	Very High
	High	Low	Low	Moderate	High	Very High
	Moderate	Very Low	Low	Moderate	Moderate	High
	Low	Very Low	Low	Low	Low	Moderate
	Very Low	Very Low	Very Low	Low	Low	**Low**
		Very Low	Low	Moderate	High	Very High

Level of Impact

We defend small shoe factories very differently than we do passenger rail switching systems – the formula does not distinguish

Eg: Spanish Rail Systems

By law…

Likelihood					
Very high	SL2	SL2	SL3	SL4	SL4
High	SL2	SL2	SL3	SL4	SL4
Moderate	SL1	SL2	SL3	SL4	SL4
Low	SL1	SL2	SL2	SL4	SL4
Very Low	SL1	SL1	SL2	SL4	SL4
	Very Low	Low	Moderate	High	Very High

Level of Impact

In practice, high-consequence assessments ignore likelihood / frequency, for unacceptable high impacts

High-End Cyber Attacks Are Not Random

Likelihood – Implies that cyber attacks are random, like random equipment failures & human errors

But … high-end attack success/failure is deterministic – the same high-end ransomware attack on the same target succeeds or fails almost always exactly as did the first one

High-end attacks use repetition – to eliminate randomness - like human errors

High-end attacks are persistent – not independent once one succeeds or when a target is strategic

The word that helps understand High Impact, Low Frequency (HILF) threats is "credible" – what is reasonable to believe – about attacks & consequences

Credibility & Frequency

Credible – defined as what is reasonable to believe (not just who)
Credible consequences – acceptable vs unacceptable
Design-Basis Threat (DBT) – most capable attack defeated reliably
What credible attacks – do we not defeat with a high degree of confidence?

Frequency				
Very High	Very Low	Low	Moderate
High	Very Low	Low	Moderate
Moderate	Very Low	Low	Moderate
Low	Very Low	Low	Low
Very Low	Very Low	Very Low	Very Low
	Very Low	Low	Moderate
		Consequence	

Out of Scope	Not Credible
Credible?	
Defeat Reliably	DBT Credible
Unacceptable	

Examples? All Past Attacks Are Credible Threats

Stuxnet – *autonomous* malware – destroyed ~1000 uranium gas centrifuges
Triton – compromised a safety-instrumented system – made mistakes though and tripped petrochemical plant, twice, before discovery
Fire in a steel mill – Israeli hacktivists started a fire that damaged equipment in an Iranian steel mill
Maersk Pharma – $1.4B NotPetya settlement in NotPetya supply chain attack
Jaguar / LandRover – $950M costs / $2B to UK economy
Colonial Pipeline & oil terminals – down for days each

Is it reasonable to believe
that these threat actors or ones like them
will target us, or others like us
within the planning horizon for our security program?

Hacktivist fire at steel mill in Iran

Reminder: Zero Days Are Credible Threats

Anthropic's Claude Mythos – is finding zero days – eg: Firefox issued an update for 320 of them

Mythos chains zero-days – taking many apparently low-severity vulnerabilities and combining them into powerful exploits

AIs are automating attacks – increasing attacker productivity 5-10x

"Boundaries should be unidirectional for the most consequential systems"

Fully-autonomous attackers will have the advantage on OT defenders

Closely track the development of autonomous attacks & invest in deterministic defenses

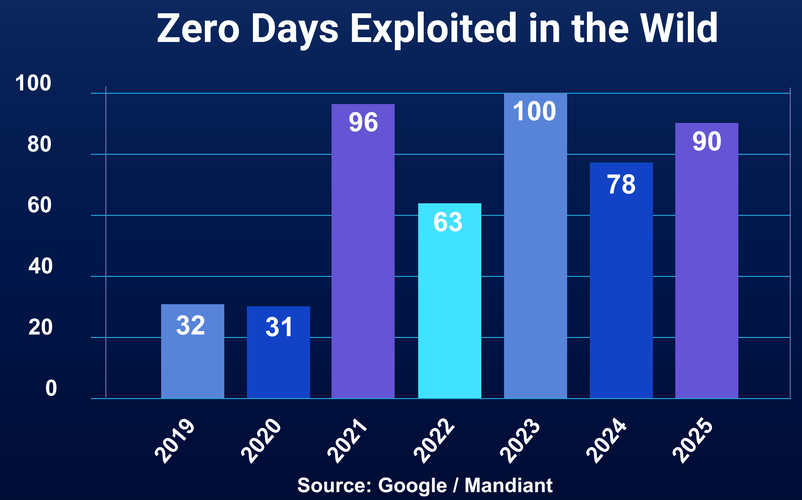

Reasonable Defences – Cyber-Informed Engineering

If your life depends on a boiler not exploding – in a cyber attack – would you prefer protection by a spring-loaded relief valve? Or longer PLC password? Where is the valve in the NIST CSF? In IEC 62443?

Manual operations – operate through compromise manually – can we still? Have we practiced? (Ukraine)

Latest – CIE Controls Database – design features that remove avenues of cyber attack or limit consequences – 62,000 records, 6000 industries, comprehensive examples of deterministic mitigations

Would you trust a switching system whose design engineer "hoped" that if an attack targeted the system, "hoped" we could detect the attack before disaster?

"Hope" is not what we expect of design engineers.

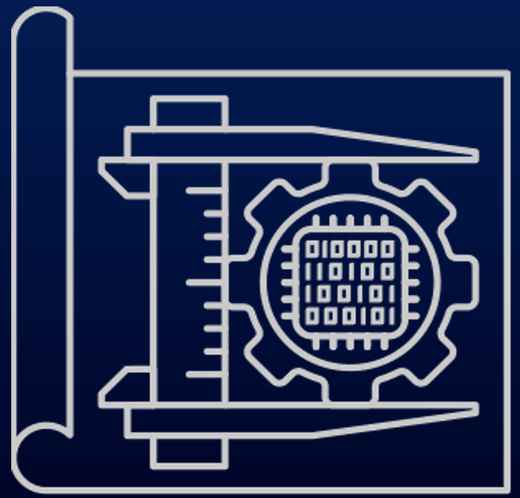

https://inl.gov/national-security/cie/

Recent – Secure Connectivity Principles for OT

Eight Authorities – UK NCSC, **Australian ASD & ACSC, NZ NCSC,** CCCS, US CISA & FBI, BSI, Dutch NCSC

Fundamental – controlling movement of attack info is fundamental, not a compensating measure

Eight Principles – balance risk / limit exposure / centralize & standardize / standard & secure protocols / harden OT boundary / limit impact / log connectivity / isolation plan

Hardware-enforced unidirectionality – to prevent OT compromise

Hardware filtering – for inbound data / attack flows

Echoes older advice from CISA / NZ – stop using VPNs, use hardware-enforced unidirectional remote access

Singapore's Code of Practice for Critical Infrastructure

OT Architecture & Security – 10.2.1 *"The CIIO shall ensure that the OT CII is not connected to any enterprise network except where necessary for operating the CII and provided that the direction of data flows is restricted to only one-way from the OT CII to the enterprise network."*

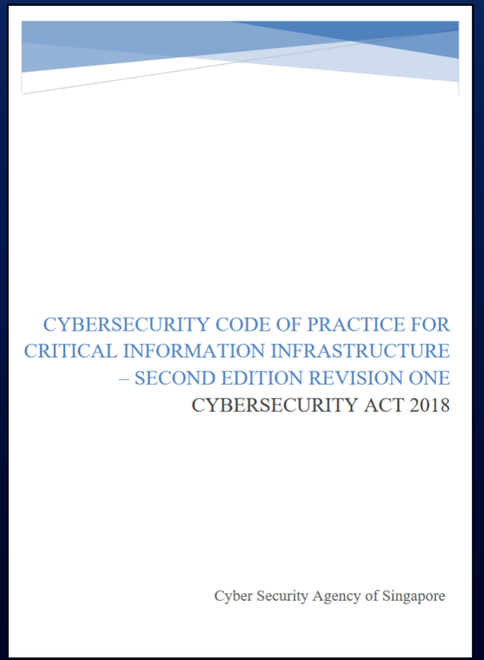

Singapore's OT Cybersecurity Competency Framework

OT Architecture & Security – 10.2.1 *"The CIIO shall ensure that the OT CII is not connected to any enterprise network except where necessary for operating the CII and provided that the direction of data flows is restricted to only one-way from the OT CII to the enterprise network."*

Network Security & Segmentation / Level 4 – *"assess feasibility for unidirectional gateway implementation in highly critical environments"*

Consistent with CIE mitigations database – hardware-enforced unidirectional gateway technology is the most universally applicable "unhackable" mitigation across all 6000 industries

Network Engineering at Consequence Boundaries

Consequence boundaries – connections between networks with very different worst-case consequences

- Safety-critical
- Reliability-critical
- Business-critical

Network engineering – engineering-grade prevention of attacks pivoting through network connections at criticality boundaries

Safety engineering shuts down ops in an emergency

Network engineering prevents emergencies from shutting down operations

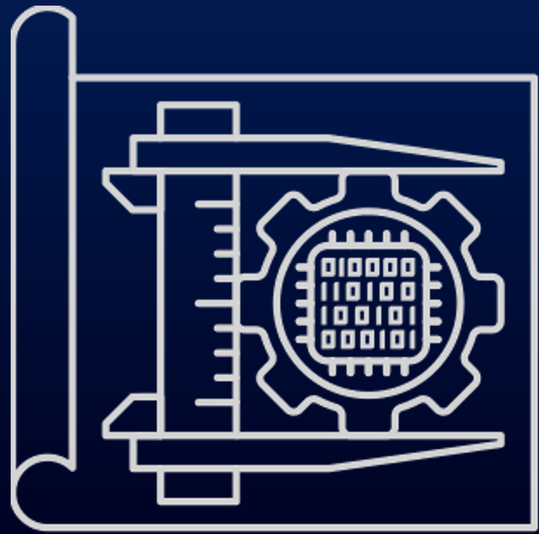

Network Engineering: Analog Signalling

Analog & binary signaling – eg: 4-20 mA current loop – impervious to malware, one-way if you use optical isolator

Works for small number of signals – gets expensive and clumsy for tens or hundreds of values that must be communicated

Deployed at consequence boundaries – most commonly to connect SIS into plant-wide automation systems without introducing cyber risk to SIS

Not possible to propagate cyber attacks into SIS using output-only analog signal paths

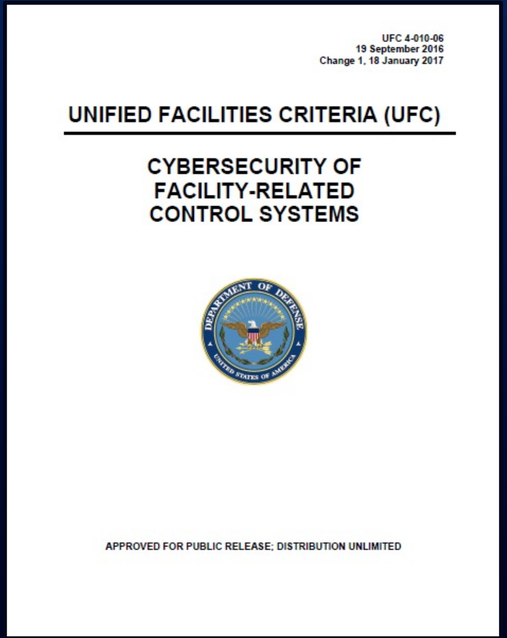

UFC 4-010-06
19 September 2016
Change 1, 18 January 2017

UNIFIED FACILITIES CRITERIA (UFC)

CYBERSECURITY OF FACILITY-RELATED CONTROL SYSTEMS

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

Network Engineering: Unidirectional Gateways

NIST 800-82 – Unidirectional gateways are a combination of hardware and software

- **Gateway hardware –** is physically able to send information in only one direction
- **Software –** copies OT servers to IT networks. IT users use the replicas normally
- **Deterministic –** cannot propagate pivoting attacks, not only “should not” or “hope not”

The only universally applicable mitigation in CIE database

Engineering-Grade Unidirectionality

Zero internal cross-connects – robust and certified unidirectional engineering

Physically divided – industrial and enterprise components

Laser / transmitter on one side + receiver / photocell on the other – not *physically* possible to send information back from receiver to transmitter

Dual power supplies – on each of sending & receiving sides

If a firewall is fully patched and correctly configured, it should not send any attacks into an OT network

With this device, even if it is not patched, or mistakenly configured, it cannot send any attacks back into the OT network

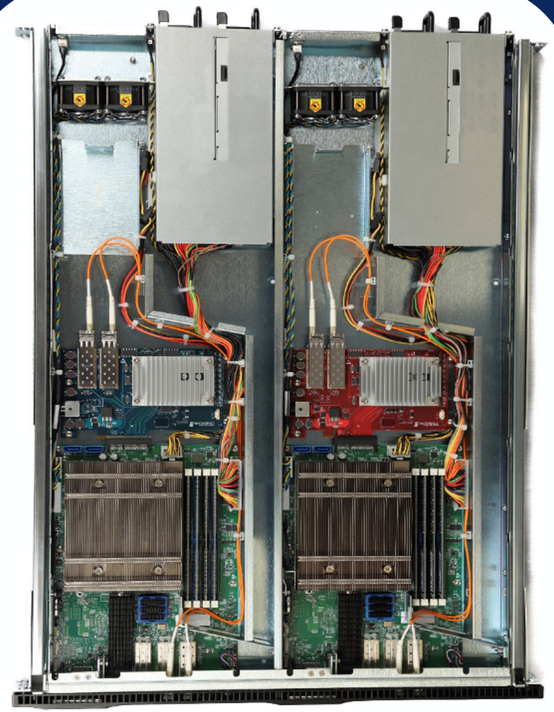

WF-600

Network Engineering: Unidirectional Gateways

NIST 800-82 – Unidirectional gateways are a combination of hardware and software

- **Gateway hardware –** is physically able to send information in only one direction
- **Software –** copies OT servers to IT networks. IT users use the replicas normally
- **Deterministic –** cannot propagate pivoting attacks, not only “should not” or “hope not”

The only universally applicable mitigation in CIE database

Credible Consequences – Blind Spots

Equipment damage – causing long-term outages – resilience demands we prevent such damage

Protective relays are software – and network-connected - equipment protection functions be defeated by simply reconfiguring the devices with much higher than safe setpoints

"Bricked" equipment – over-write firmware so that devices cannot be booted up far enough to restore good firmware – brick 80% of out-of-support PLCs in a plant or wind farm – do we have enough spares to come back?

Supply chain – when geopolitical tensions turns into kinetic conflicts, will we continue to trust software updates from "the other" side of the conflict, or their allies? What will we do with our systems sourced from those nations?

These capabilities have been demonstrated in the wild – credible threats

Acceptable Risk – What Do Insurers Think?

Lloyds regulator – £200M+ coverage for cyber damages requires special review, nation-state exclusion, dropped silent coverage

Due care insurance questionnaires – increased from less than one page to more than 5 pages of questions, including questions about unidirectional protections

Large businesses self-insure – for risks that Lloyds won't touch? Is that reasonable? Does JLR regret self-insuring?

When we accept risk – the business should charge the division accepting that risk with what would be the cost of an insurance policy, had we purchased one?

$5B in general damages & general liability cover – including cyber consequences – was available 10 years ago – not today

Today other insurers followed Lloyds lead

LLOYD'S

What Is Reasonable?

Protect against widely-available capabilities – intent can change in a heartbeat – literally

Nation-state ransomware = pervasive threat – demands network engineering for unacceptable consequences / optional for acceptable consequences

Compliance – essential, but judge will ask “did you do what was reasonable?” Not “did you do the minimum the law required?”

Network engineering to address ransomware threat + strong NIST / IEC 62443 posture to manage insider threat

How Much Is Enough?

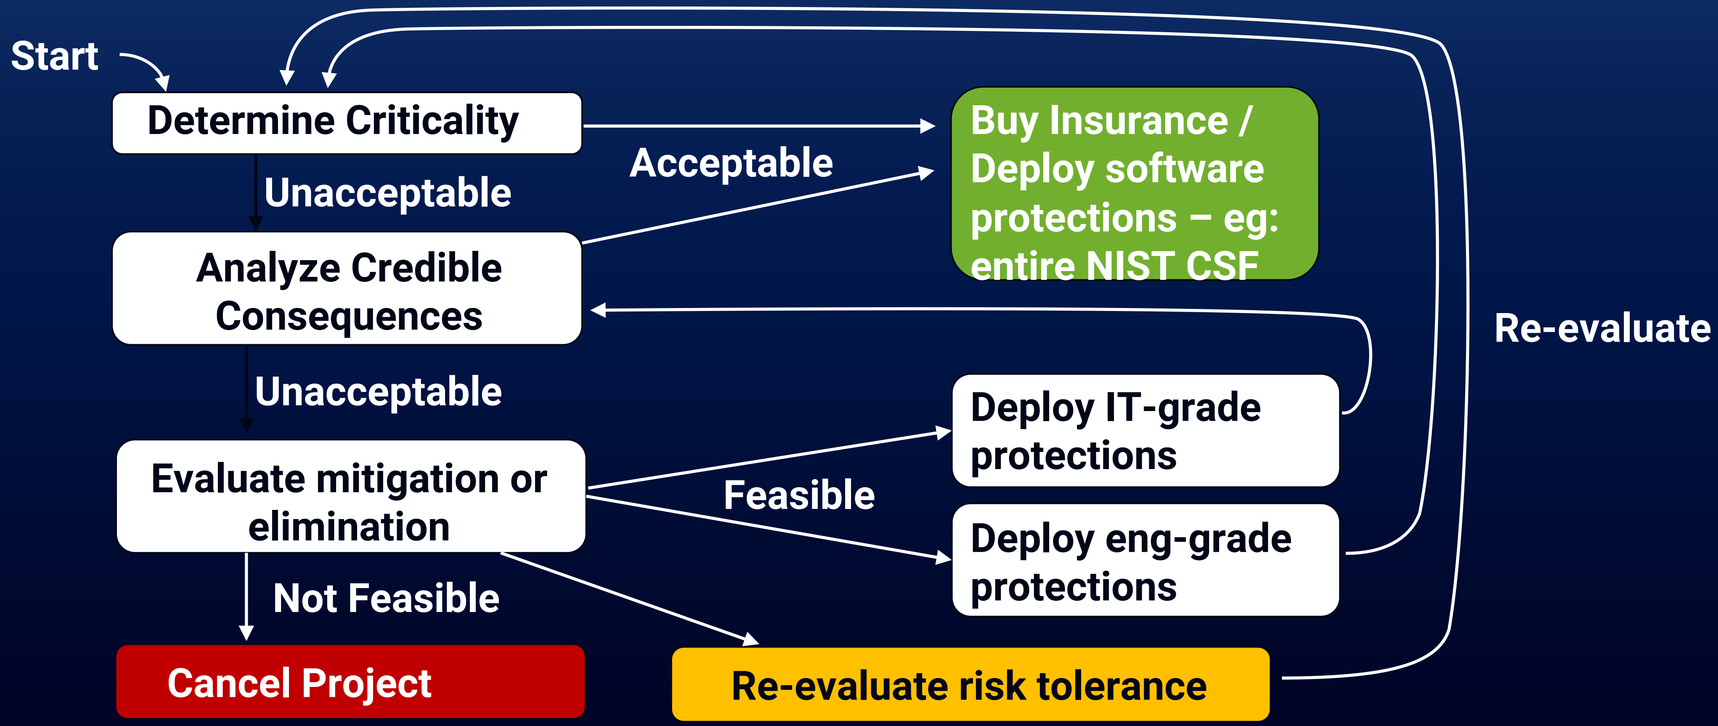

Reasonable Responses to Credible Threats

Unacceptable consequences – generally must be defeated with a high degree of confidence – ie: engineering / safety-grade - IT-grade is often not enough

Mitigation cost / effectiveness – engineering-grade mitigations are almost always more effective and often cheaper than IEC 62443 SL4

Return on security investment – engineering-grade protection is enormously cheaper than the worst outages – eg: JLR ($1-2.5B) and Maersk Pharma ($1.4B)

Credible consequences – not "James Bond" scenarios

Re-evaluate risk tolerance – executive or board should sign off on new policy when material risks are accepted

"As strong as is reasonably practicable" – is the stated or implied requirement in many jurisdictions

Practicable = feasible

Reasonable = what any other reasonable person in the role would decide – not just what you think

Hacktivist fire at steel mill in Iran

Example – Compromised SIS (Triton)

Credible consequence – pivot through IT network into OT, reaching engineering workstation – reprogram SIS & operate PLCs to cause unsafe condition SIS no longer stops explosion

IT-grade defenses – can reduce probability of detect / respond – eg: to 97% ??

97% is not enough – safety consequences are expected to be defeated to safety standards, not IT standards

Network engineering – at IT/OT interface (eg: unidirectional), or SIS + safety workstation / DCS interface (eg: analog signaling) – defeats pivoting with a high degree of confidence

Photo credit: Robert Fries <license>

How Much Is Enough?

Study attacks – to understand credible threats

Defend information flows – they are all attack vectors

Deterministic defenses – are “future-proof” – analog, electro-mechanical & deterministic digital designs

Also deploy NIST CSF / 62443 / etc. – for completeness & to address “easier” LIHF risks

Communicate residual / accepted risk – in short messages to execs or BoD – ideally have them sign off

When safety is at risk – or critical infrastructure, or large amounts of shareholder money – deploy defenses that are “as strong as is reasonably feasible”

Deploy engineering-grade *“future proof” analog, electro-mechanical & deterministic digital defenses when credible consequences are unacceptable*