From CS 491/691 Active Cyber
October 28, 2026

Demo of generative AI Perplexity, and the Turla cyber-actor

Hi! I am a Computer Science professor, and teaching a course in cybersecurity. Can you tell me about a threat called TURLA?

Hi, Charles! **Turla** is a long-running, Russian state-linked cyberespionage group—not simply the name of one virus. MITRE tracks it as G0010 and attributes it to Russia's Federal Security Service (FSB). Its documented targets include governments, embassies, military organizations, universities, research institutions, and pharmaceutical companies in more than 50 countries, with activity dating to at least 2004.

For a cybersecurity course, Turla is a useful case study in how an intelligence operation combines targeted intrusion, custom malware, and covert communications rather than relying on a single exploit.

What the name means

There is some naming confusion because reporting sometimes uses related names for both the actors and their tools:

Name	**What it refers to**
Turla	The threat group or associated espionage activity; MITRE identifier G0010.
Waterbug, Venomous Bear, KRYPTON	Alternative names used in reporting for the actor.
Snake / Uroburos	Names associated with sophisticated malware used in Turla operations; reporting sometimes overlaps these names with the group itself.
Carbon	A separate, modular backdoor framework used by Turla.

The useful classroom distinction is: the actor conducts campaigns, campaigns use techniques, and techniques may be implemented through several different malware families. Turla should not be treated as one executable with one fixed set of capabilities.

What Turla wants—and how it operates

Its central mission is espionage: obtaining sensitive information from strategically valuable organizations. That makes it different from a financially motivated ransomware operation whose most visible objective is to force a payment. Reporting on Turla emphasizes intelligence collection and carefully selected targets.

ESET's analysis of the Carbon framework describes a particularly instructive, staged intrusion pattern:

1. Initial access. A target receives a spearphishing message or visits a compromised website that the target normally trusts—a "watering hole."
2. First-stage foothold. A relatively lightweight backdoor, such as Tavdig or Skipper, provides initial access.
3. Reconnaissance. The operators examine the compromised environment before deciding where to deploy more sophisticated tools.
4. Selective deployment. Carbon is installed on systems considered worth further attention.
5. Command and control. Separate components load the malware, dispatch tasks, communicate with operator-controlled infrastructure, and interact with other computers on the network.

This is a documented example, not a universal sequence for every Turla intrusion. Its teaching value is the deliberate separation between "get in" and "commit valuable capabilities to this target." ESET also documented repeated Carbon revisions, illustrating that the group changes its tools rather than maintaining a static malware inventory.

Why Snake became especially notable

Snake used compromised systems as more than places from which to steal information. Infected machines also participated in a peer-to-peer network that relayed communications, helping conceal the relationship between the operators and their ultimate targets. Its communications used a custom encrypted protocol, complicating detection and investigation.

That creates an important distinction for students:

- A compromised computer can be an intelligence target.
- It can also be infrastructure used to conceal attacks against somebody else.
- The same computer can serve both roles.

In May 2023, the United States announced Operation MEDUSA, an operation against the Snake network. The FBI developed a tool called PERSEUS that issued commands causing Snake to overwrite vital components and disable itself. U.S. authorities publicly linked the operation behind Snake to a unit within the FSB's Center 16.

That was a disruption of a particular malware network—not evidence that every Turla capability, operator, or intrusion had disappeared. The actor–tool distinction matters here too.

How I would frame it for your students

A compact case-study exercise could use MITRE's Turla profile alongside ESET's Carbon analysis. The former provides the actor-level view; the latter gives students a concrete intrusion sequence and software architecture to examine.

Suggested discussion questions:

- Attribution: What evidence would connect a malware sample to a campaign, a campaign to an operator, and an operator to a government? Which of those conclusions requires the most caution?
- Architecture: Why separate an initial foothold from a more elaborate second-stage backdoor?
- Detection: What evidence might remain when malware runs within a legitimate process or relays traffic through another victim?
- Incident response: Why might removing one known malware family be insufficient to establish that an intrusion is fully remediated?
- Measurement: Does "number of infected hosts" measure intelligence value, infrastructure scale, or both?

Those are proposed teaching questions rather than claims about every Turla operation. They lead naturally from the documented staging in Carbon and the relay architecture in Snake.

For an introductory lecture, I would summarize the case this way: Turla illustrates an enduring espionage actor whose identity is broader than any one malware family—and

whose operations make the relationships among initial access, persistence, collection, and covert communications especially worth studying.

What are Turla's key intrusion techniques

How do Turla and Agent BTZ connect

What are good student exercises on Turla

How does Turla use satellite C2 networks

What are notable Turla detection strategies