Security and Performance

Paul Kummer
Head, e-Infrastructure and IS Security Officer
CCLRC

Overview

- **So – what is the problem?**
  - **Data rates**
  - **Hacking**
  - **Viruses and SPAM**
  - **The Web**
- Some general comments on security
- Towards a solution

“Typical” traffic levels (1)

- RAL – 1 minute resolution

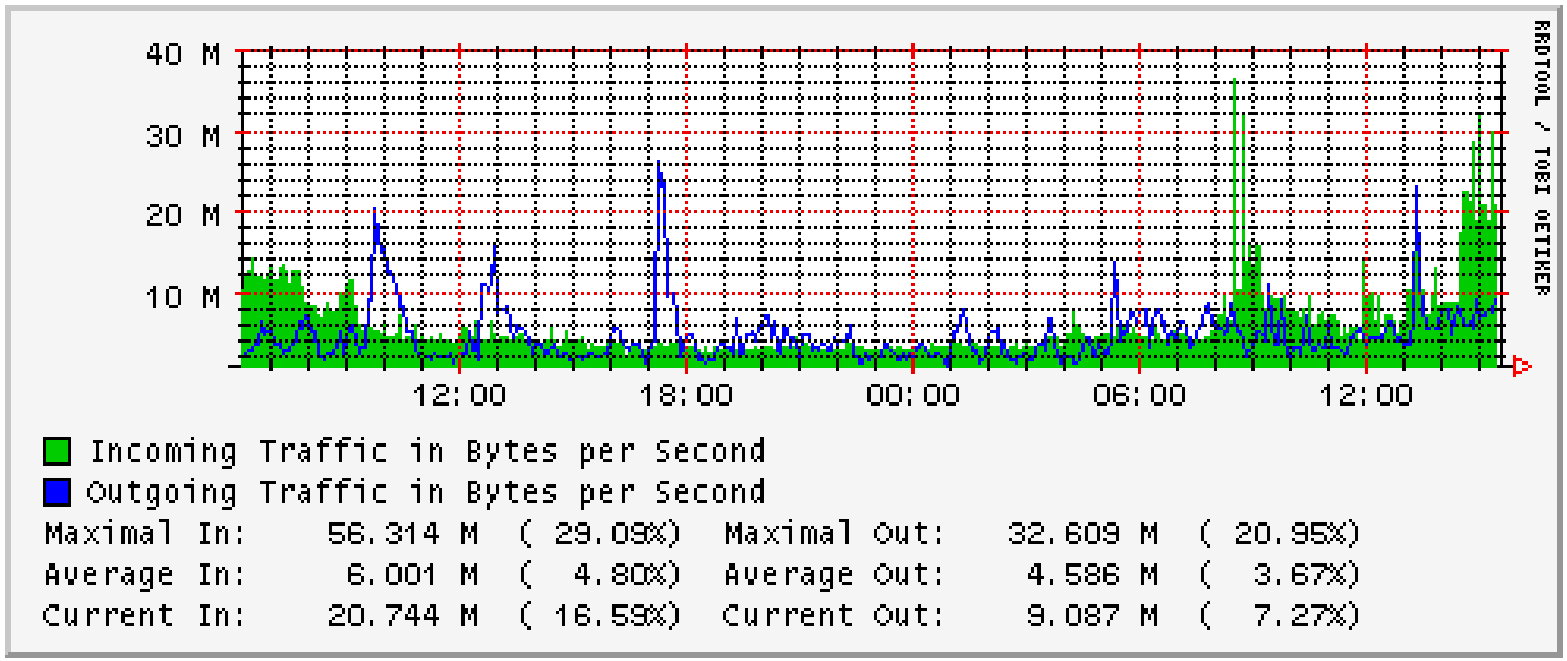

“Typical” traffic levels (2)

- RAL – 1 day resolution

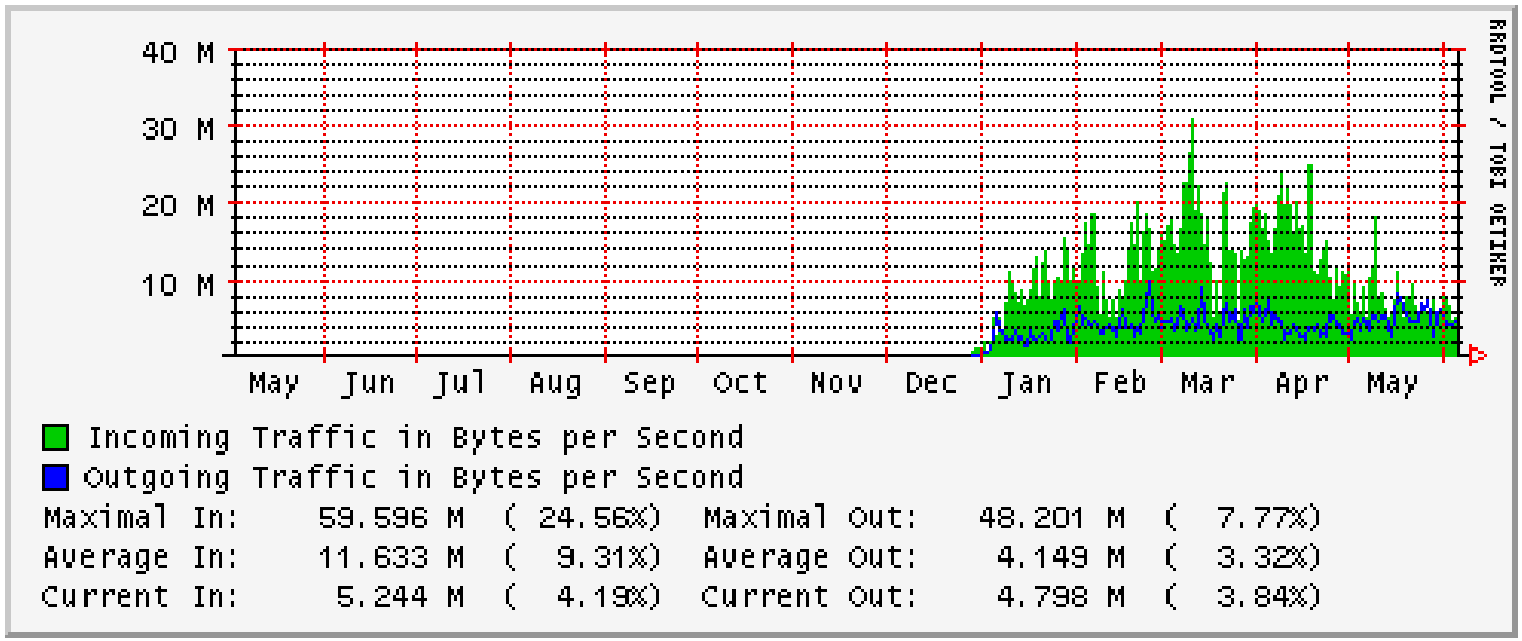

LHC – service challenge 2

- CERN - >600MB/s daily average for 10 days
  - ~5Gb/s

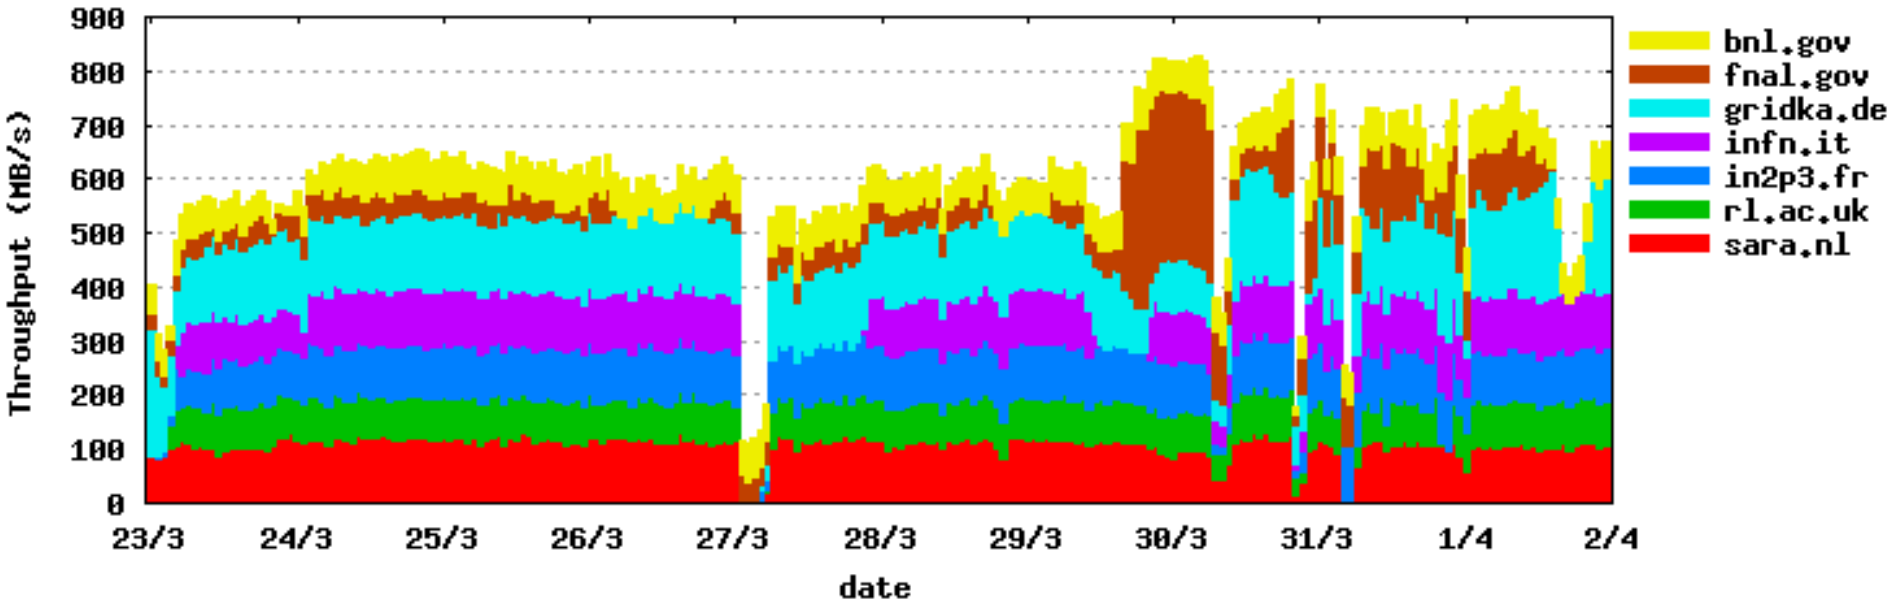

Hacking probes

- Each CCLRC site receives about 30,000,000 probes a day looking for a weakness in the defenses.
  - 300/second
  - Firewall log is about 5GB/day (uncompressed)
- Average compromise time now measured in days
  - Successful probe → compromise can be measured in seconds

Do viruses and SPAM matter?

- Depends on bandwidth
  - ADSL can be totally compromised by peer-to-peer file sharing
  - 1Gb/s link is unlikely to be affected by email fluctuations

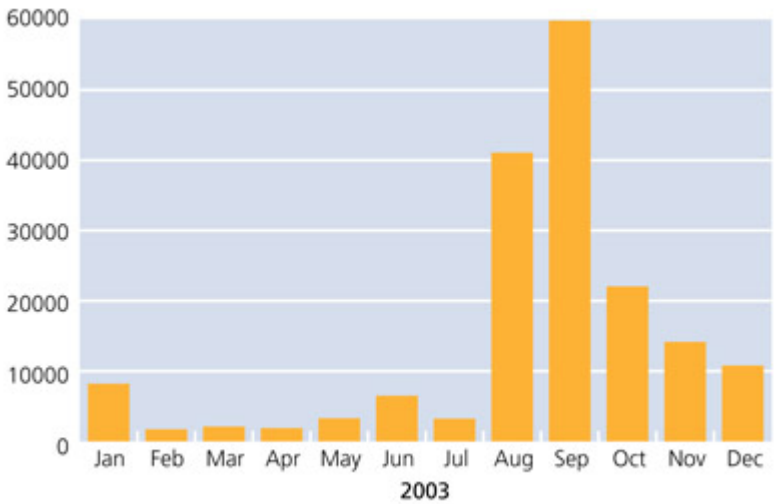

SoBig-F

1600 viruses/day
@100kB each
160Mbytes/day
15kb/s (average)

Email and SPAM (1)

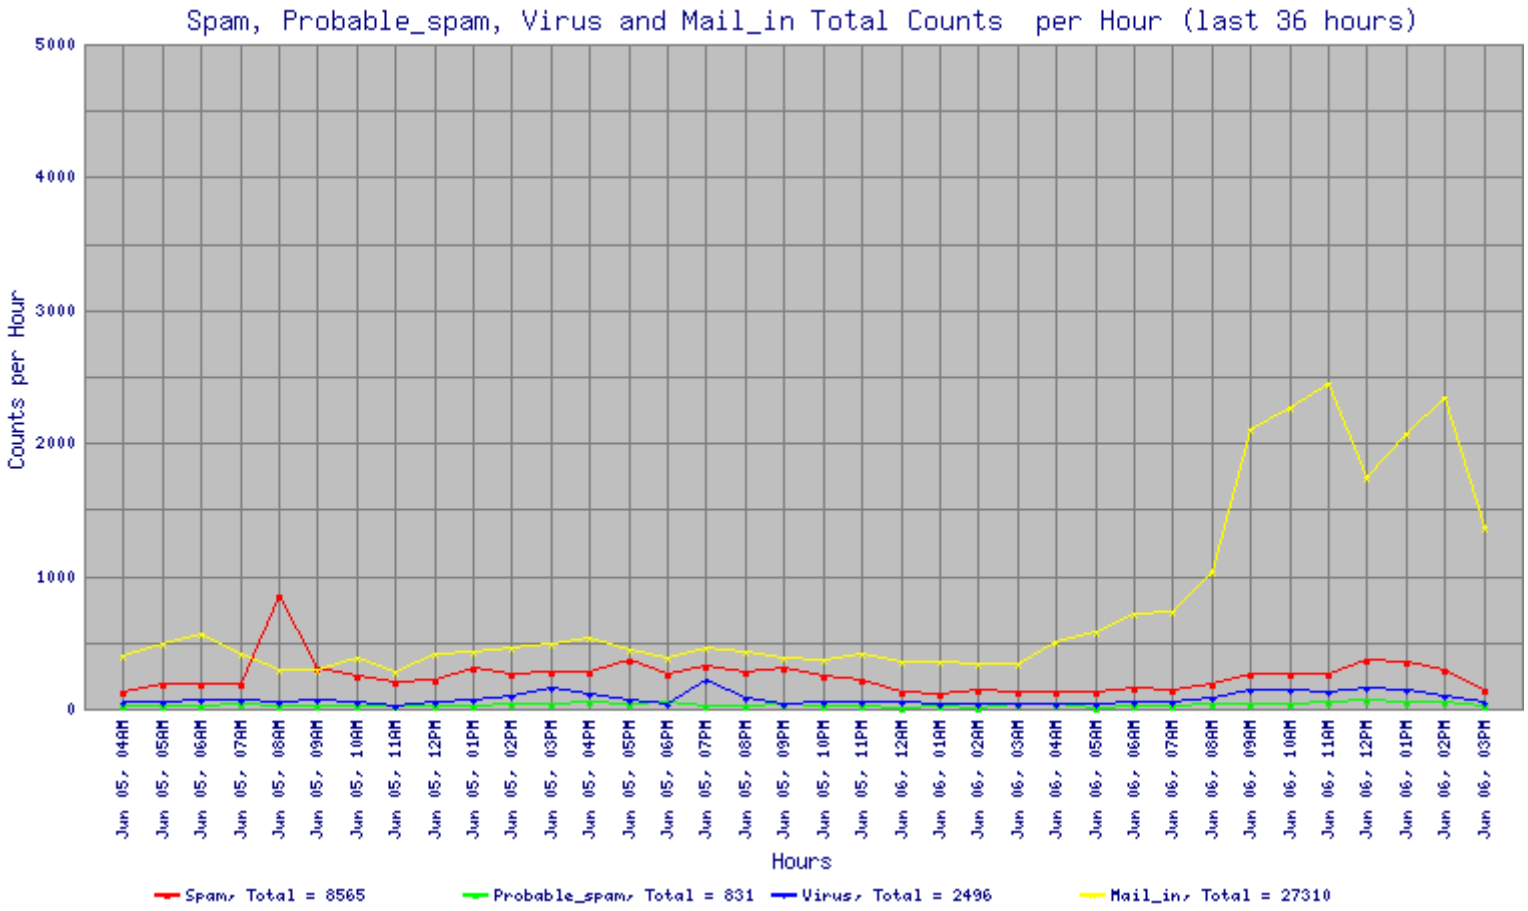

Email and SPAM (2)

- SPAM
  - 500/hr @ 10kB
  - 11kb/s
- Email
  - 2500/hr @ 10kB
  - 55kb/s

And then there is the Web

- 500 people at Daresbury Laboratory generate about 2Mb/s averaged over the working day.
- Traffic is bursty
  - 1-100 connection setup/cleardown per second

Overview

- So – what is the problem?
- **Some general comments on security**
  - **Risk analysis**
- Towards a solution

Risk analysis (1)

- Scientific environment usually needs more “flexibility” than a commercial environment
  - “Unusual” protocols
  - “Need” to “do your own thing”
    - Fewer controls over the individual
- Can never get absolute security
  - The “enemy” is dynamic
  - Constant need to keep protection up-to-date
    - Currently measured in hours for viruses

Risk analysis (2)

- May need to trade security against bandwidth (against cost)
- Bandwidth for LHC > capability of current firewalls
  - And gigabit firewalls (if suitable) are expensive
- Security in depth
  - Multiple layers

Overview

- So – what is the problem?
- Some general comments on security
- **Towards a solution**
  - **Structure**
  - **Firewalls**
  - **Access control lists**
  - **End-system tools**
  - **The Grid**
  - **Certificates and encryption**

“Standard” security structure

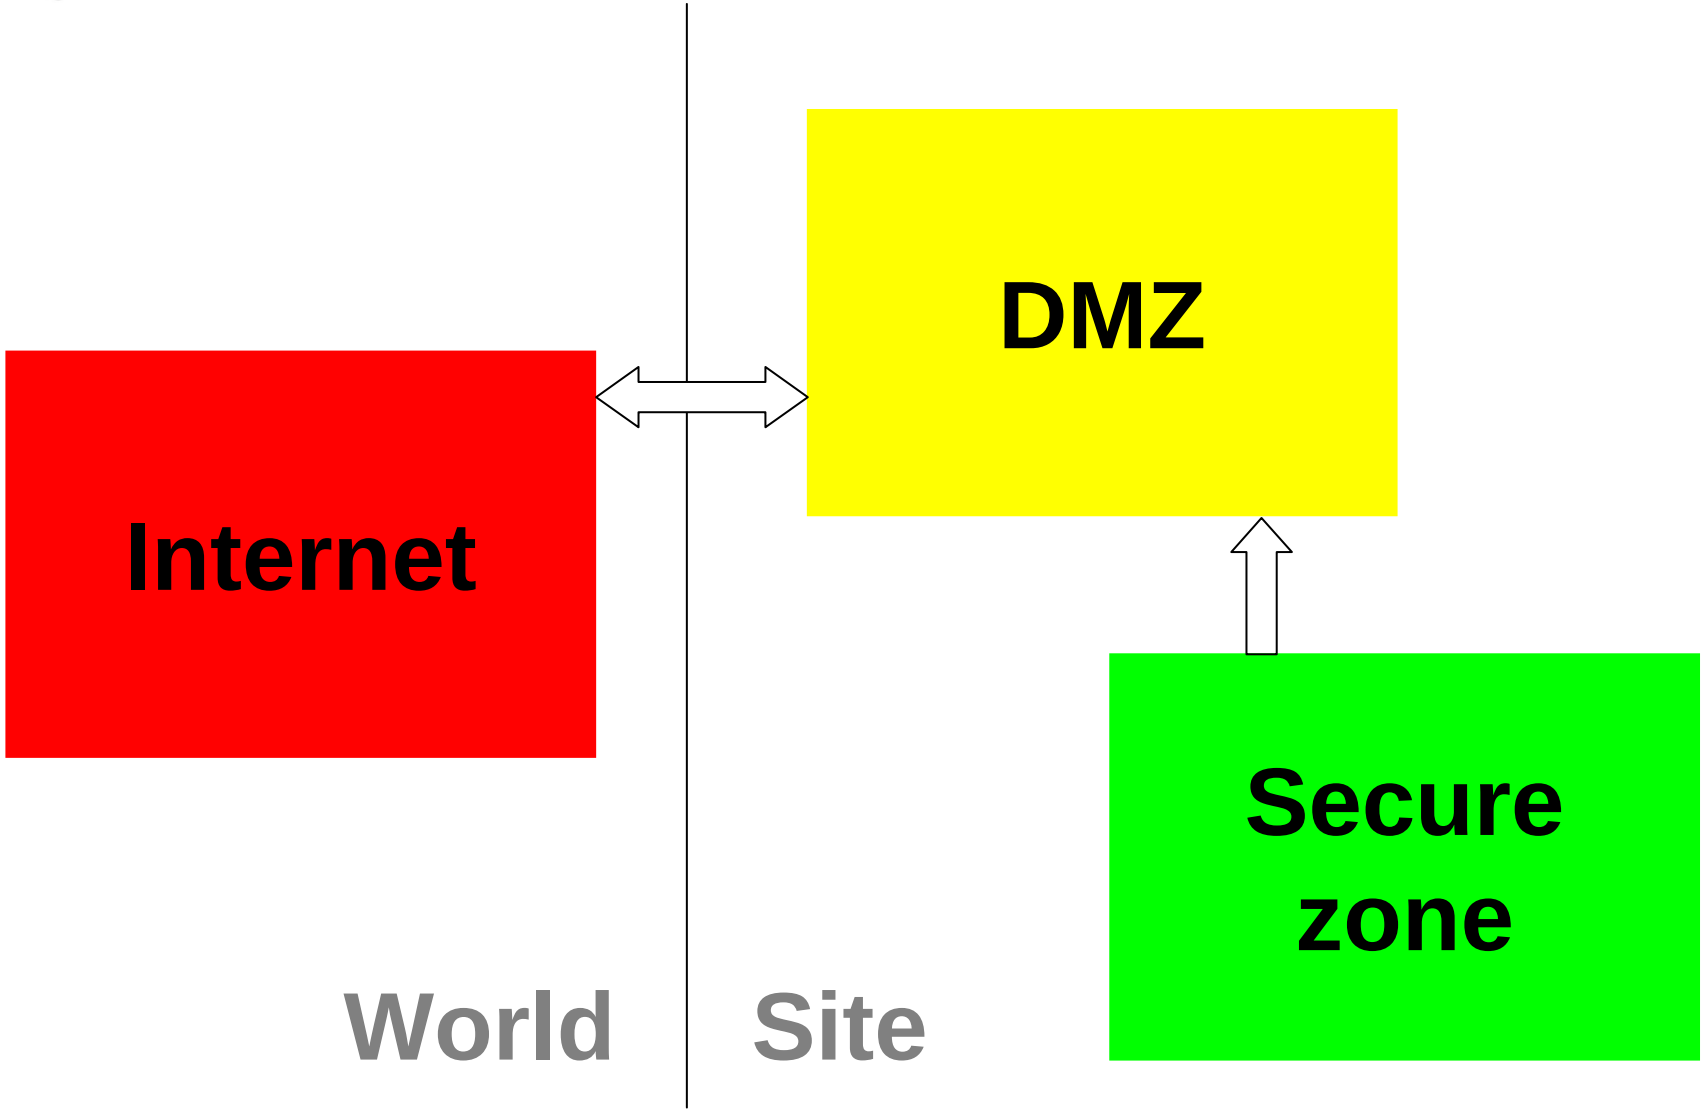

CCLRC model

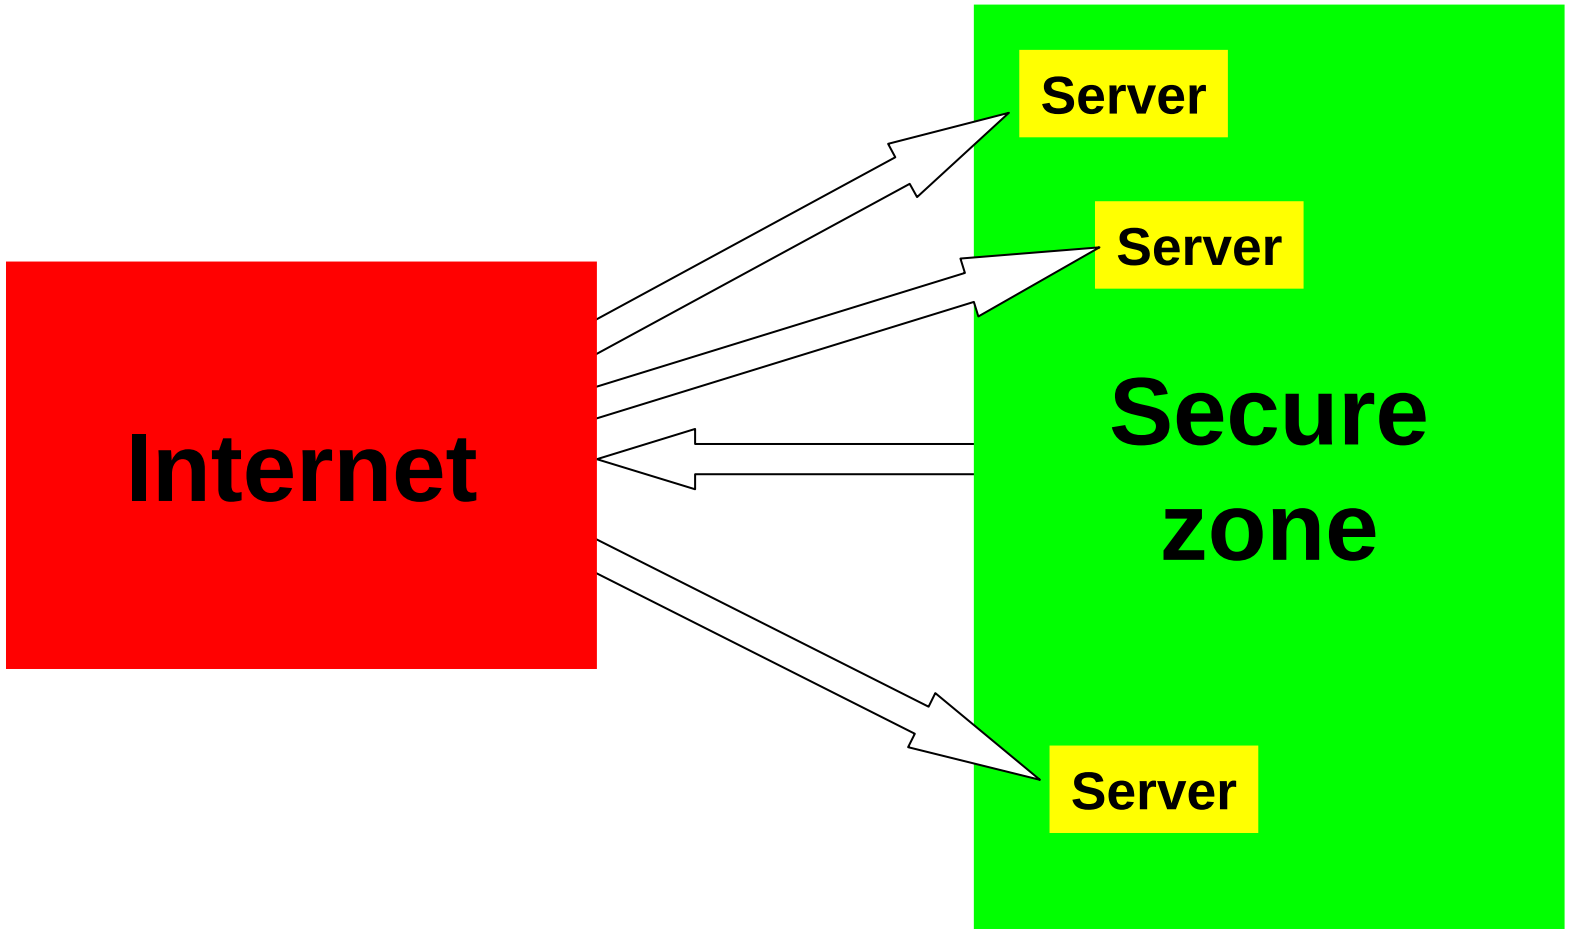

Firewall (1)

- Keep “state” for each communications session
- Interpret the data stream to get state
- Policies used to accept/deny communications
- Detect and stop DoS attacks
- Detect port and address scanning
- Potential performance bottleneck

Firewall (2)

- Bottleneck prevention
  - Buy a firewall based on processing capability – not link speed.
  - Special purpose hardware

1Gb/s Ethernet interfaces:	8
Concurrent sessions:	1,000,000
New sessions/second:	25,000
Firewall performance:	up to 4 Gbps
Triple-DES (168 bit) performance:	up to 2 Gbps
Policies:	40,000
Rules:	200,000

Firewall (3)

- Note
  - Maximum throughput < total link speed
  - Maximum throughput degrades if 3DES used
  - Higher session startup per second → better DoS resilience
  - UDP communications count towards session count
  - Session information includes:

    Source (IP address : port) : Destination (IP address : port)

Firewall (4)

- Firewalls handle “problem” protocols
  - E.g. FTP

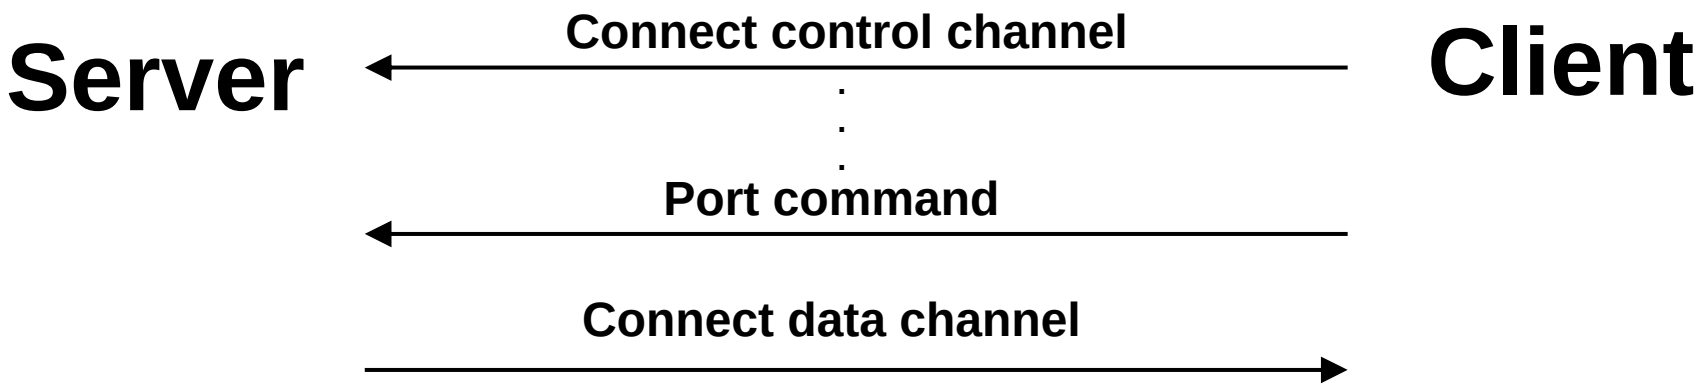

Firewall (5)

- Firewalls do not handle “special”, problem protocols
  - Multi-stream FTP where several data channels are opened to get extra throughput
    - GridFTP
    - BBFTP
- Don’t expect commercial firewalls to recognise the latest protocols

Firewall (6)

- The broadcast problem (also applies to switches)
  - Broadcast frames need to go out on multiple ports
  - May be handled by the control processor (especially in chassis-based systems)
  - The control processor is much slower than the special purpose hardware
  - May be a bottleneck

Access control lists (1)

- Not necessarily state based
- Control restricted compared to a firewall
- Usually based on TCP/IP and UDP/IP information
  Source (IP address : port) : Destination (IP address : port)
  TCP flags
  - The latter is used to distinguish connect requests from all subsequent packets

- Typically:

```
Src=Any,Dst=148.79.242.4:80 Allow
Established Allow
```

Access control lists (2)

- Disadvantages compared to firewalls
  - No DoS protection
  - Cannot handle “problem” protocols

```
Src=Any:FTPdata,Dst=Any:1025-65535 Allow
```

- Advantages compared to firewalls
  - Often available in large switches (low cost)
  - Much higher performance (line rate)

Access control lists (3)

- Useful in a controlled environment
  - A limited number of systems

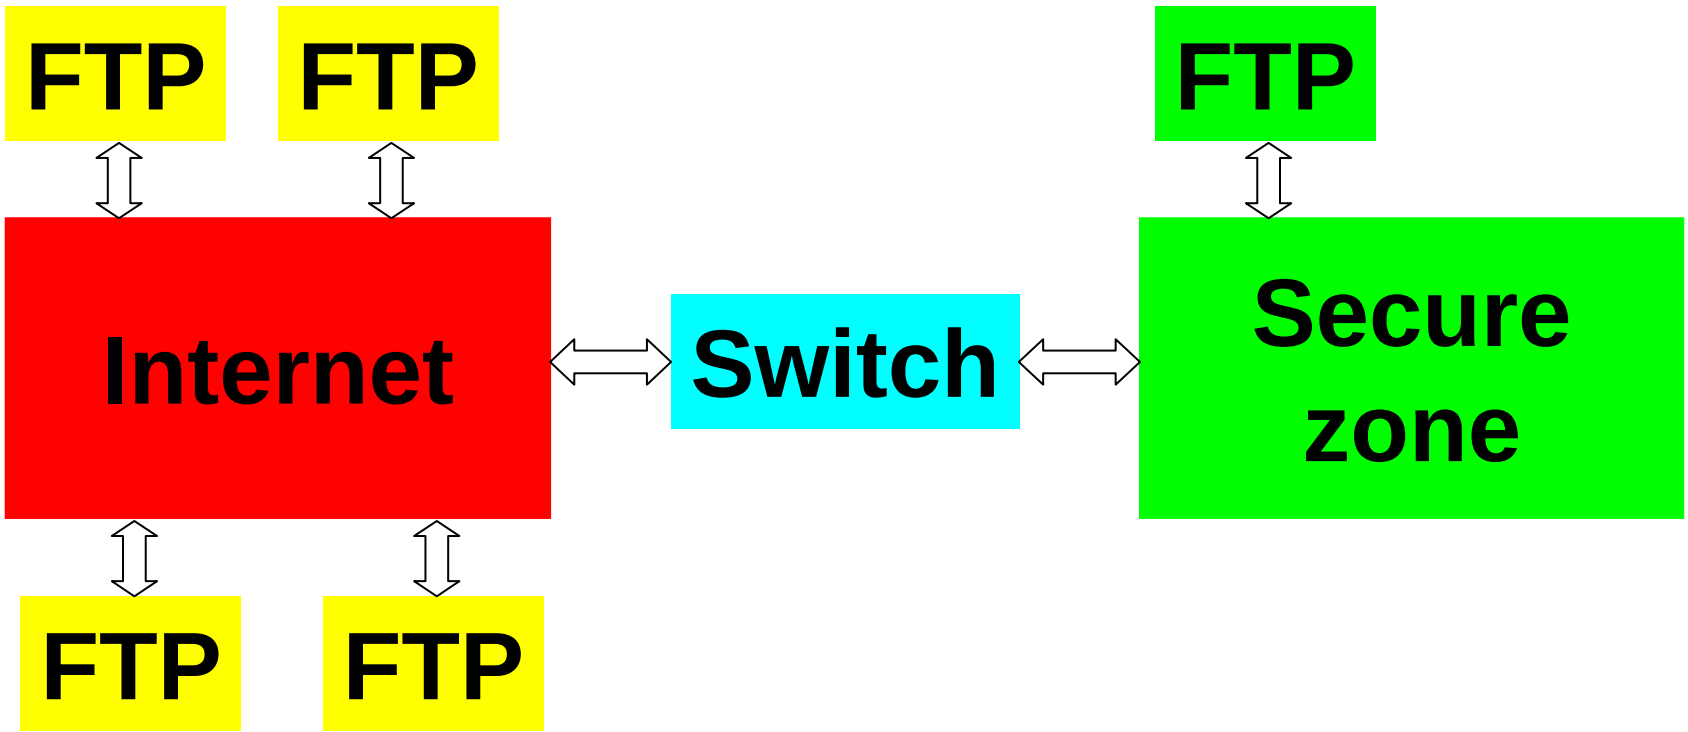

Access control lists (4)

- Could combine ACLs with Firewall

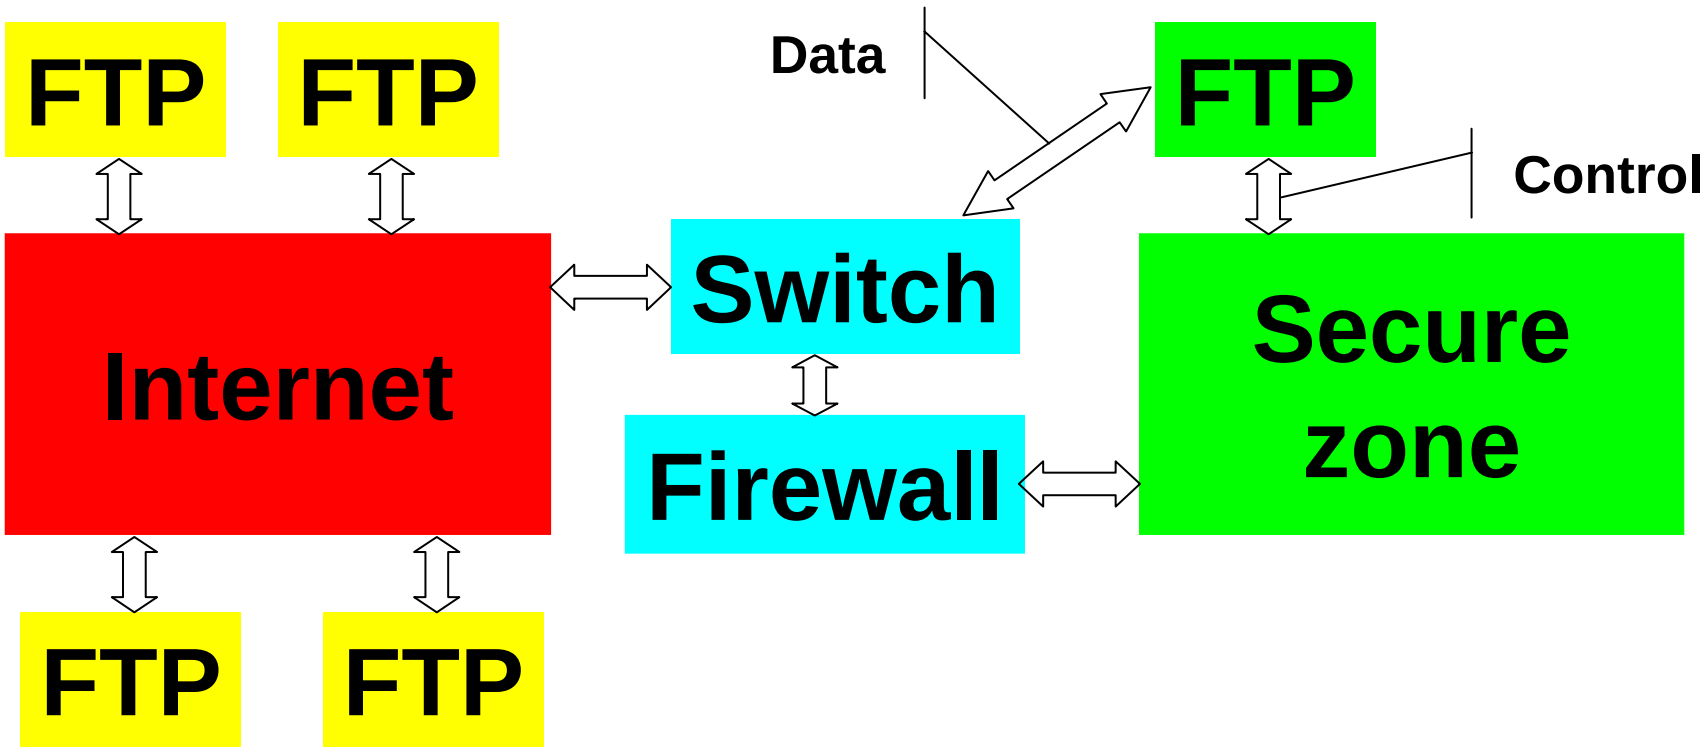

End-system tools (1)

- Linux
  - IPchains / IPtables
    - Both are packet based
- Windows
  - Personal firewall (many)
    - Packet based
  - Anti-virus (many)
    - Byte based (examines the data stream)

End-system tools (2)

Byte based

- Virus checking becomes feasible below 100Mb/s
- Special purpose hardware gives 100-1000Mbit/s throughput

2GHz processor

Bandwidth	Instructions/byte
10Gb/s	2
1Gb/s	20
100Mb/s	200

End-system tools (3)

Packet based

- Affect on throughput is dependent on packet size
  - NOT the TCP buffer size
  - BUT the IP packet size
    - Subject to reduction all along the communications path
    - Typically 1500B on LAN
    - Can reduce to 256B on WAN
    - Note “big frames” on Ethernet (8kB)

The GRID (1)

- GRID security is based on certificates
  - High level of security between systems
  - Implies high level of trust
  - Takes no account of low-level attacks
    - E.g. buffer overruns

The GRID (2)

- Design is not “firewall friendly”
  - GLOBUS - requires multiple ports to be opened
    - System ports (≤ 1024) + range above 1024
  - Web services likely to be worse
  - (Almost) reduces a firewall to a switch with ACLs
- Web services on port 80 a problem
  - Default may go through web cache
  - Managing “exceptions” may not be scalable

Certificates and encryption (1)

- Cryptographic techniques operate on byte streams
- Performance dependent on:
  - Encryption type
  - Hardware/software implementation
  - Operating system (I/O, memory management)
  - API and its implementation
- The network may not be the bottleneck

Certificates and encryption (2)

- Simplified application protocol

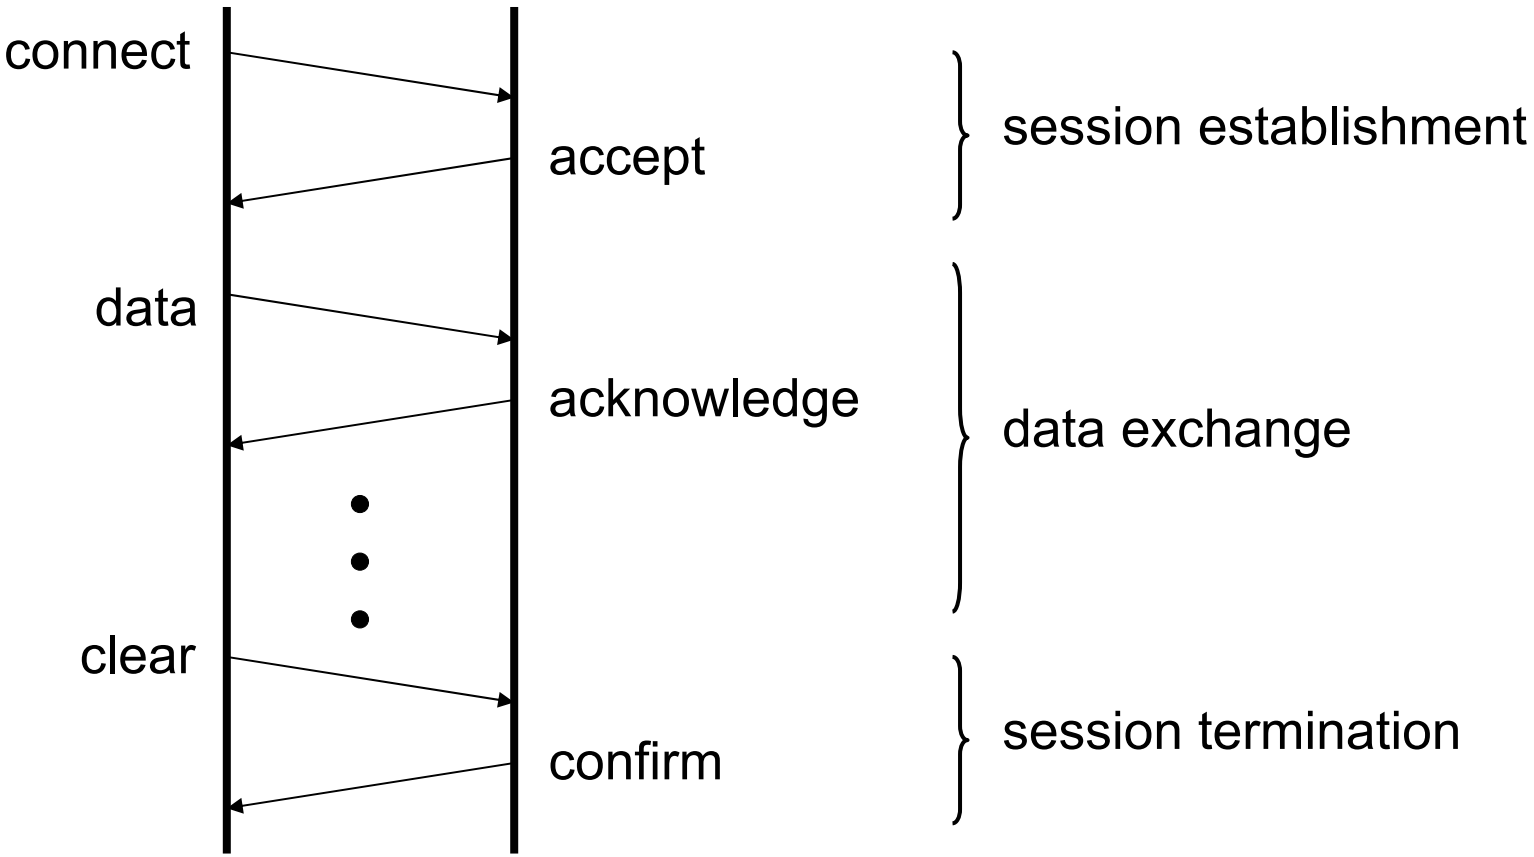

Certificates and encryption (3)

- Certificates exchanged during session establishment

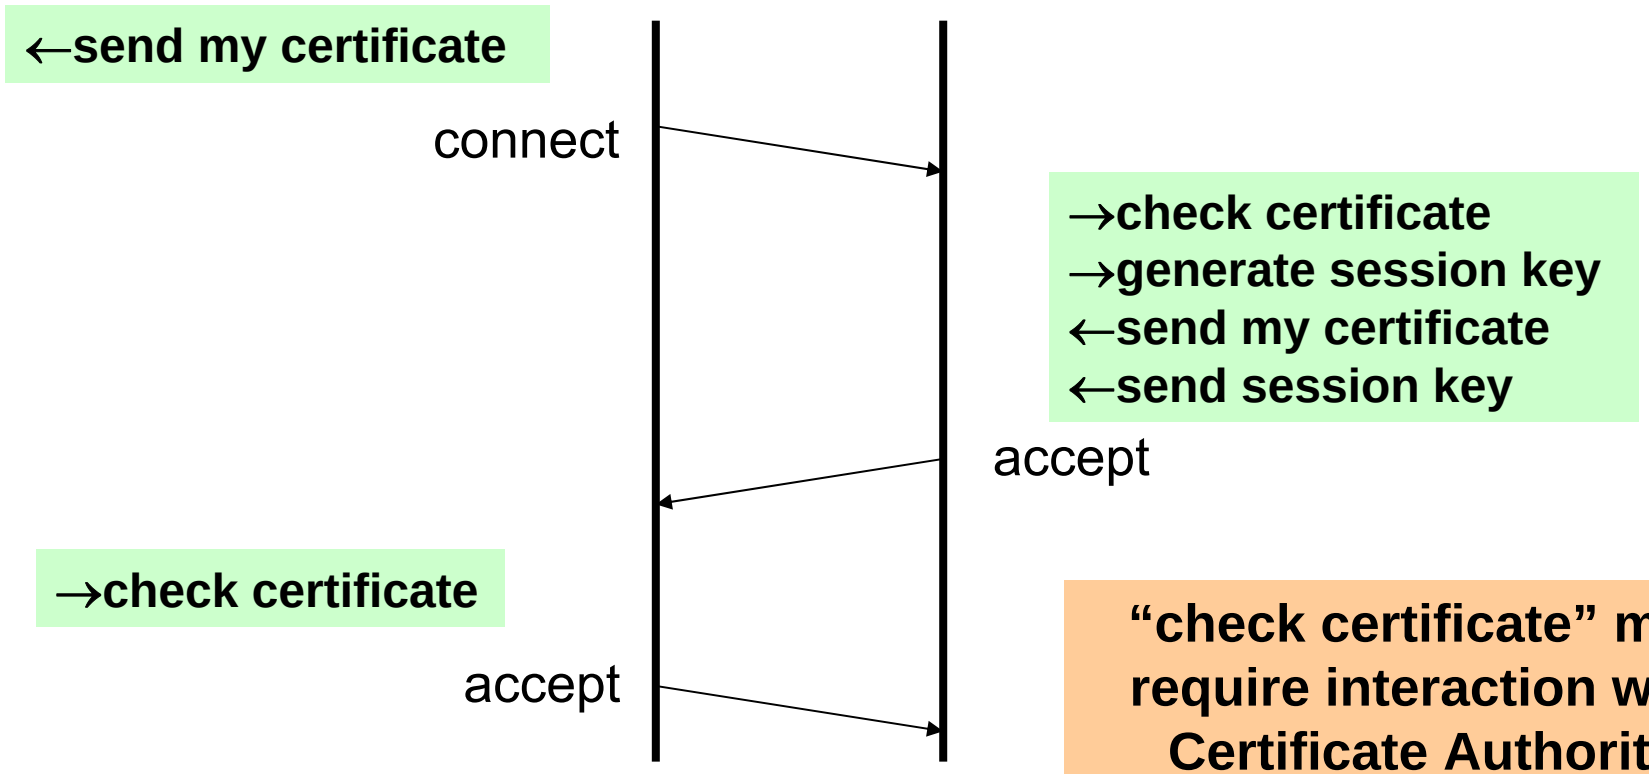

Certificates and encryption (4)

- Session key used to encrypt data