



System and Service Quality: a return to the good, old values?


*Giorgio Ventre
Università di Napoli Federico II
Laboratorio ITEM del Consorzio CINI
giorgio.ventre@unina.it*

- Open issues in QoS (sort of...)
- From QoS to RoS (Reliability of Service)
- Neptune: a network emulation project

- In spite of all hype on global network architectures, today we face a complex and heterogeneous reality:
 - ✓ Fixed access networks: POTS, xDSL, CATV, MetroLAN
 - ✓ Mobile, wireless access networks: GPRS, UMTS, WiFi
 - ✓ Interoperability with terrestrial digital broadcasting
- Additional complexity issues:
 - ✓ New, diverse terminals: (Symbian Cell. Phones, PDAs, smart set-top-boxes)
 - ✓ Dynamic creation of novel services and applications

- The availability of a multiplicity of networks, devices and services should be seen as an opportunity:
 - ✓ No single infostructure of critical importance
 - ✓ Ease of access to all players: government, companies, common people
 - ✓ Availability of a multitude of sources of information
- ... provided that such rich scenario can be managed as a system.

- Most of the existing networks are separately and statically capable of hosting different applications/services offering differentiated services
 - ✓ Service Level Agreements (SLA) : the user view
 - ✓ Service Level Specifications (SLS) : the system view
- In an highly heterogenous and mobile scenario we probably need a new metaphore: the Service Condition



- Open issues in QoS (sort of...)
- From QoS to RoS (Reliability of Service)
- Neptune: a network emulation project


- We have got some lessons recently:
 - ✓ 9/11 2001 Attacks
 - ✓ US East Coast Blackout
 - ✓ Italy Blackout
 - ✓ Series of attacks:
 - Worms (NIMDA, ...)
 - D-DOS
 - Routing attacks
- We probably need to re-discover traditional values typical of traditional engineering practice

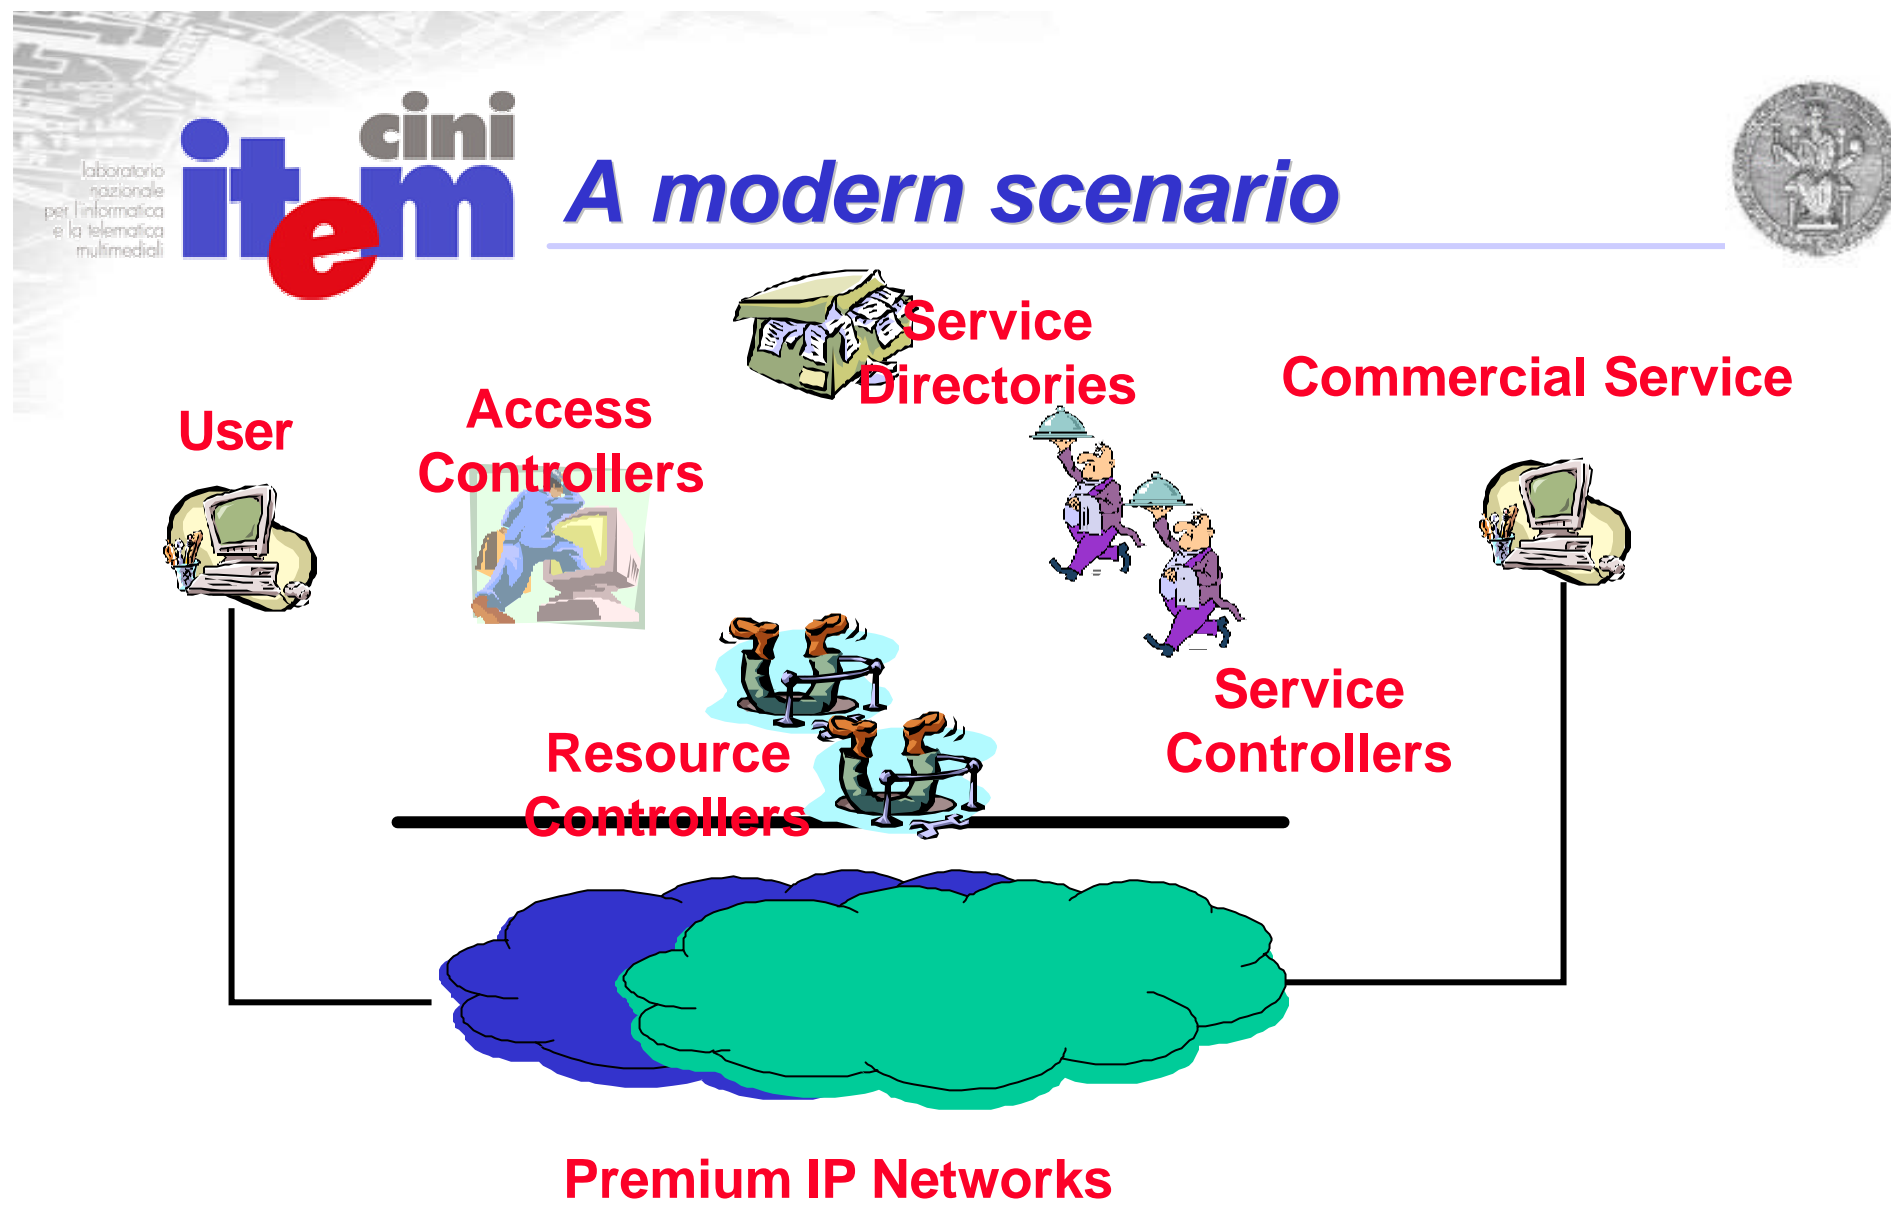
- “The Internet under Crisis Conditions” A Committee of the National Research Council of the National Academies (www.nap.edu):
- Findings of the Committee
 - ✓ Attacks had very limited effects on the Internet as a global, best effort communication system
 - ✓ Effects were much harder on some specific locations and services
 - ✓ Internet technology appears to be robust *per se* but considerable efforts are needed to protect internet-based systems
 - ✓ Many critical interdependency discovered only after the attacks

- Dependency of Internet on other TLC systems (fixed, wireless, cellular)
 - ✓ Obvious: co-location of sites, tubes, cables
 - ✓ Not so obvious : e.g. communications between NYC ISPs and TelCos hampered by problems to toll-free numbers
- Facility disaster planning as a rare expertise/culture in the Internet world
 - ✓ Very limited capacity of backup power generation even in major ISP sites/POPs
- Clear abuse of Internet flexibility
 - ✓ DNS for .za domain was hosted on a server in NYC
 - ✓ WiFi LANs of two major Manhattan hospitals operating in outsourcing via Internet

- Anticipated by the US East Coast blackout: much larger scale than WTC but apparently more limited damages
- Different effects and impacts
 - ✓ POTS infostructure capable of enduring very long power outages: practically no effects
 - ✓ Cellular Networks locally in deep crisis
 - ✓ National TV and Radio broadcasters Ok, local players generally in crisis
 - ✓ “Global” and VoIP operators knocked-out
- What about the Internet?
 - ✓ All IT based services affected : AAA, CDN, Servers

Networks today are not only wires





- Forget OSI layering/abstractions
 - ✓ Services depend not only on peer and adjacent layers
 - Resiliency is a system-wide issue, with vertical and horizontal dependencies
 - Start speaking about networked systems and not only of networks
 - IT based services must be considered as part of the whole picture
- Contributions from several disciplines

- Different issues for different scenarios
 - ✓ Robustness of services
 - To unexpected situations: faults, excess in demand, soft attacks (D-DOS)
 - To expected but complex situations: tools/methodologies for proper dimensioning of services (Service Engineering)
 - ✓ Resiliency of infostructures
 - Focus on survivability of communication systems to hard attacks (terrorist hits, natural disasters)
 - ✓ Reconfigurability of communication systems
 - Make different networks as a single infostructure

- Open issues in QoS (sort of...)
- From QoS to RoS (Reliability of Service)
- Neptune: a network emulation project

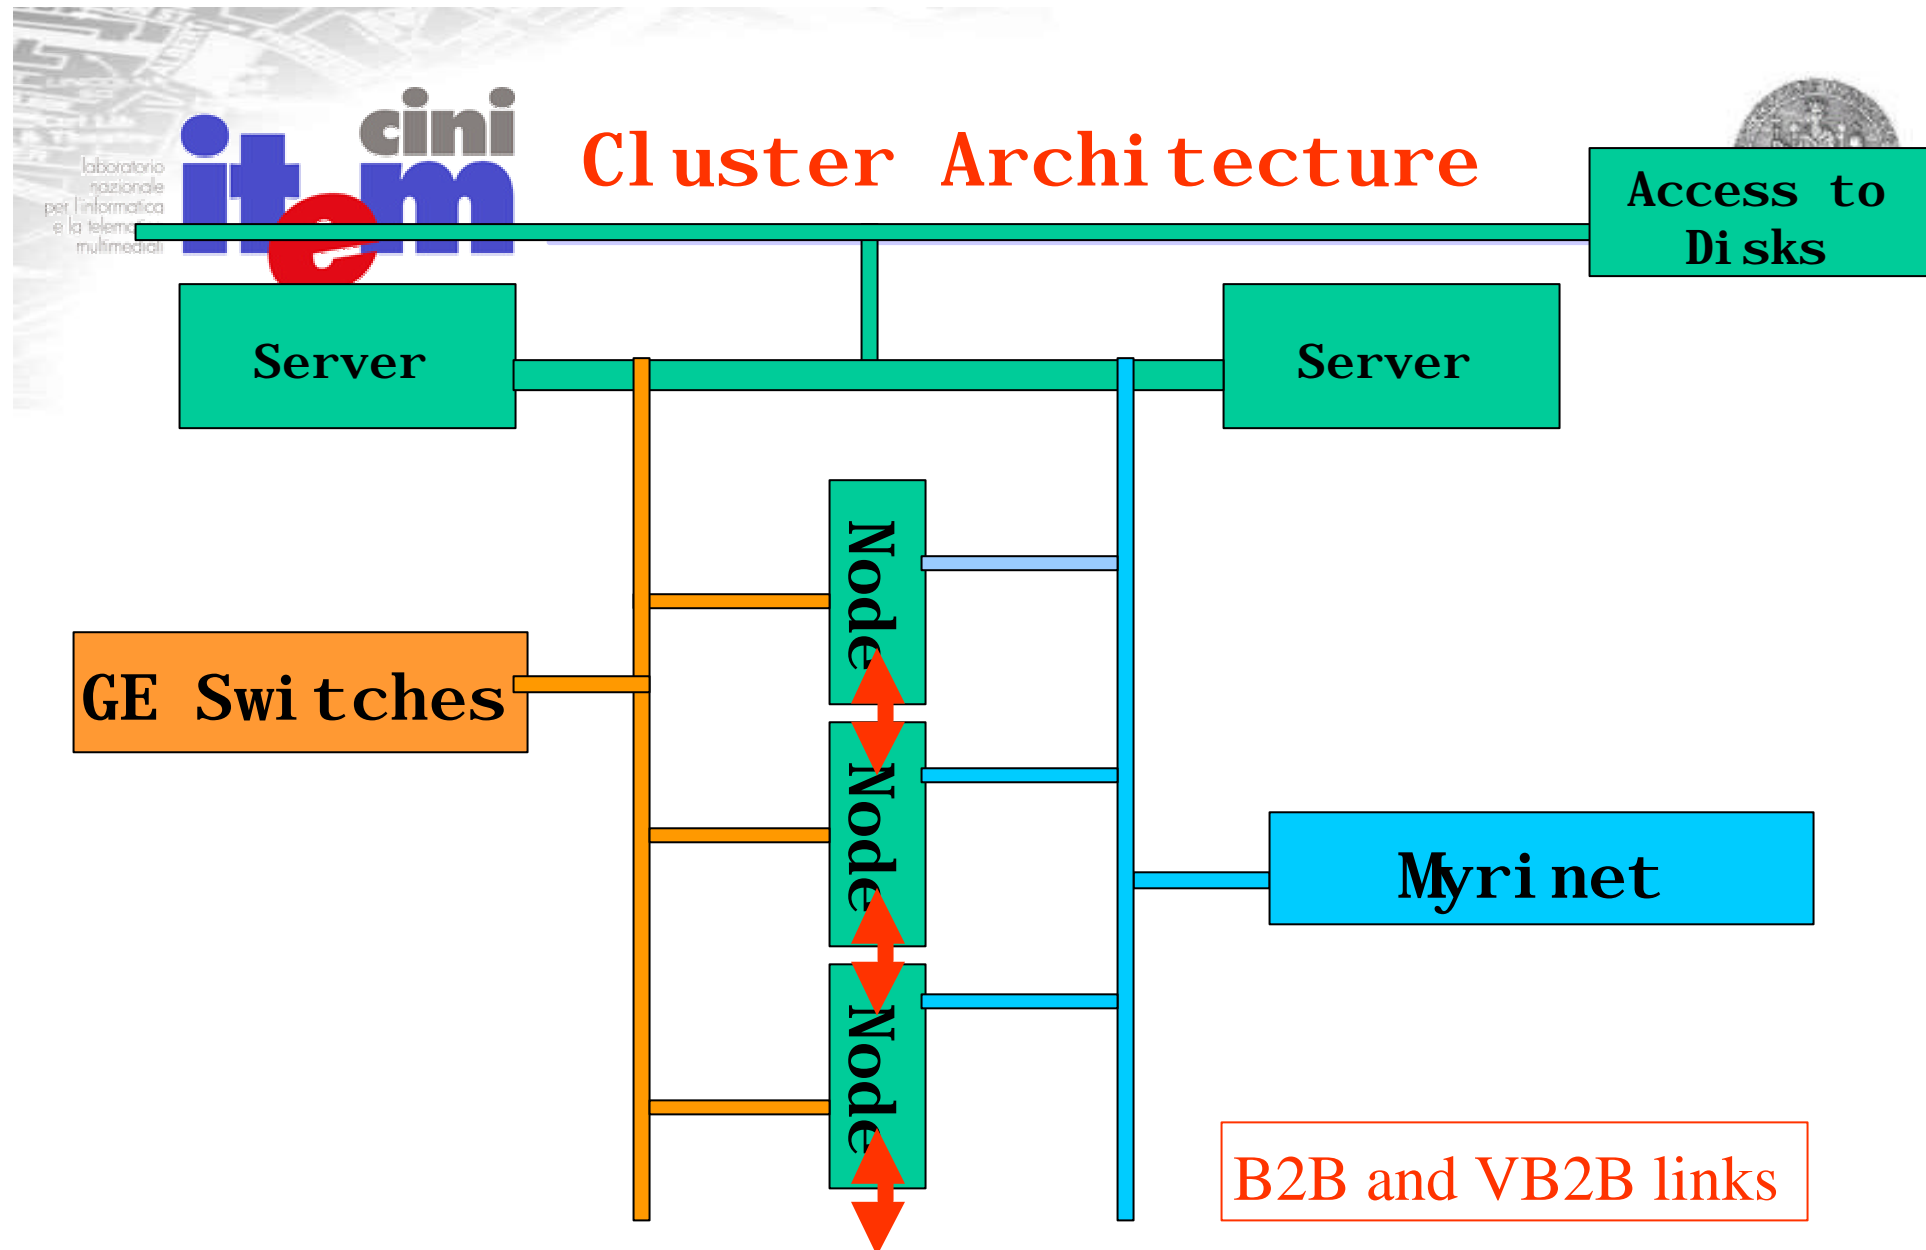
- How complex is today's Internet?
 - ✓ Large topologies
 - ✓ Highly variable service scenarios
 - ✓ Heterogeneous edge and core entities
 - ✓ Stocastical source behaviour and distribution
 - ✓ Complex traffic patterns
 - ✓ Proprietary forwarding gear
 - ✓ Uncoordinated deployment of novel services and devices

- Can simulation help us?:
 - ✓ Small scale simulations, ok!
 - micro-scale behaviour
 - High-level, functional analysis
 - ✓ Large scale simulations, nok!
 - micro-scale simulations computational intensive for realistic topologies
 - Macro-scale evaluations highly depending on phenomena dynamics
 - Hybrid approaches difficult to realise
 - ✓ Generally we simulate protocols but we ignore how they are actually implemented in terminals and nodes
 - ✓ Need to consider computational aspects for applications, nodes, systems

- Computing nodes can emulate network elements
 - ✓ Terminals
 - ✓ Applications & traffic sources
 - ✓ Complex distributed systems
 - ✓ ... Network nodes
- However, large scale network emulations demand for large scale computing power
- Is this feasible?

- Network Emulation for Protocol TUNing and Evaluation
- Use cluster computing technology to make feasible large scale network emulations
 - ✓ Large number of computing nodes
 - ✓ Focus not only on computing power but also on node connectivity: multi-grids for complex topologies
 - ✓ Exploit Open Source technology for programmability and flexibility
 - ✓ Interworking with proprietary systems
- Develop more powerful techniques and tools for emulation control and management.

- A rack-based system
 - ✓ 2 bi-processor management nodes (to RAID, peripherals)
 - ✓ 30 bi-processor computing nodes
 - Hyperthreading technology
 - Local disk
 - ✓ Multiple communication interfaces per node
 - 2 GE (switch based access to/from management nodes)
 - 4 FE (switch based and back-to-back interconnection)
 - 1 Myrinet interface for out-of-band, high-speed signalling and monitoring data collection
 - PCI slots for future extensions
- Connectivity to proprietary systems (switches/routers)



- The cost of Monitoring
 - ✓ Monitoring and communication of monitoring data should not interfere with emulation processing and data flows
- Hyperthreading for reduced monitoring overhead
 - ✓ 1 processor as 2 different processors:
 - 1 node as emulation engine
 - 1 node as monitoring probe
- Multiple, specialised communication interfaces
 - ✓ Myrinet infrastructure dedicated to:
 - Fast data gathering towards RAID
 - Signalling to-from management nodes