


Data networks and storage networks

(some of the things you wanted to find out but never got around to trying them to see what really happens)

Colin Butcher



Part 1:

- Basic principles of data networks and storage networks

Part 2:


- Data network protocols (TCPIP, DECnet, SCS, LAT, MOP, AMDS etc.)

Part 3:

- Network infrastructures - putting it all together


Part 4:

- Examples & discussion




Part 1:

- Basic principles:
 - Data networks - LAN & WAN
 - Storage networks - SAN





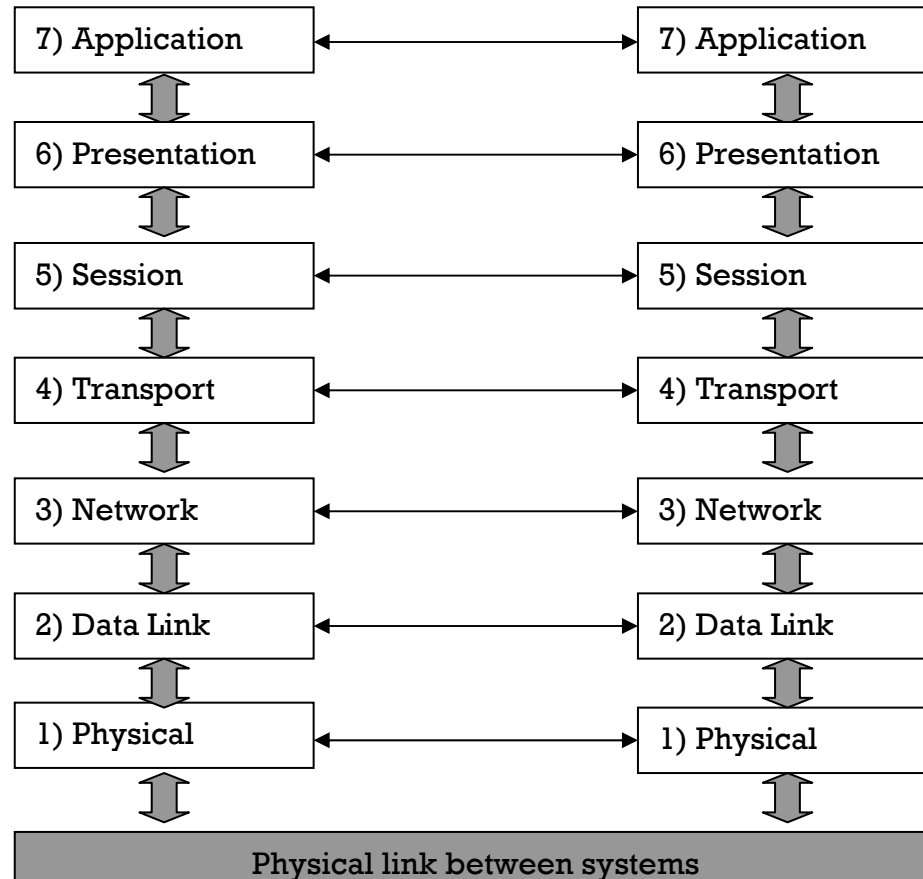
Data networks

- Ethernet technologies
- Physical components and cabling
- Protocols and addressing
- Network Interfaces
- Network Segmentation
- Wide-Area networks (WANs)




7	Application	Provides for distributed processing and access, contains application programs and supporting protocols (eg FTAM)
6	Presentation	Coordinates conversion of data and data formats to meet the needs of the individual applications
5	Session	Organises and structures the interactions between pairs of communicating applications
4	Transport	Provides reliable transparent transfer of data between end systems with error recover and flow control
3	Network	Permits communication between network entities
2	Data link	Specifies the technique for moving data along network links between defined points on the network, and how to detect and correct errors in the Physical layer (layer 1)
1	Physical	Connects systems to the physical communications media







- Transmission properties are important - a square wave at in one end needs to be recognisable as a square wave coming out at the other end
- Copper:
 - Co-axial (thick-wire, thin-wire)
 - Twisted pair (Category 5, 5E, Category 6 etc.)
- Fibre-optic:
 - Monomode (typically 9 micron)
 - Multimode (typically 50 or 62.5 micron)




- 10 Mbit/sec
 - 100 Mbit/sec (Fast ethernet)
 - 1,000 Mbit/sec (Gigabit ethernet)
 - 10,000 Mbit/sec (10Gigabit ethernet)
 - Copper / fibre (different transmission characteristics)
-
- FDDI, Token ring




- Wireless ethernet (2Mbit / 11Mbit / 54Mbit / etc.)
 - Note: GPRS, HSCSD, HSDPA, Bluetooth etc. are not wireless ethernet
 - Access points connect wired and wireless LANs
 - Shared bandwidth
 - Complex to configure and manage in a large environment
 - Access control and data privacy are major issues



- Provide connection between IO subsystem and network
- Copper / fibre / wireless physical interfaces
- On-NIC processing:
 - Packet creation
 - Address filtering
 - Encryption
 - Protocol processing (TCP/IP offload - TOE)



- Hardware MAC address
- Physical MAC address
- Broadcast address
- Multicast addresses
- Point to point addresses
- Ethernet packet format v IEEE802.3 packet format
- Packet size (normal frames and jumbo frames)
- Jumbo frames - frame size varies with NIC type and must be supported by every device all the way through the infrastructure




Why segment a network?

- Performance
- Security
- Availability


How can you segment a network?

- Multiple NICs
- Repeaters
- Bridges
- Switches
- VLANs
- Routers




- Layer 1 devices (“flat” network)
- Provide electrical fault isolation
- Simply re-time and re-transmit signal
- No control of bandwidth
- Beware of cumulative end to end delay exceeding maximum permissible frame timing – which leads to ‘folklore’ such as the “three repeater rule”


- *TIP: Beware of the generic term “hub”*



- Packet content based (Layer 2)
- Store and Forward
- Easy to use and configure
- Poor control of bandwidth (filtering)
- Spanning tree algorithm
- Provides an extended LAN
- Not all protocols can tolerate the inherent delays in working over an extended LAN
- Remote booting (MOP, BOOTP etc.) will absorb bandwidth




- Introduces parallelism
- Speed of chipsets (latency & bandwidth)
- Full duplex operation on a single device per port basis
- Traffic monitoring (mirror ports)
- Link aggregation
- Bandwidth control
- “Store and forward” versus “Cut through” switching
- Layer 2, Layer 3 etc. switching
 - Layer 3 generally refers to TCP/IP routing layer
 - Layer 4 generally refers to TCP/IP port numbers, eg: HTTP port 80 traffic)




VLANs are another way to segment a network for performance and security


- Implemented within core switches
- Also implemented in some NICs / device drivers
- Port based VLANs
- Protocol based VLANs
- Connectivity between VLANs
- VLAN tagging of packets (802.1Q)
- VLAN tagging of packets out of NICs
- QoS (Quality of Service) and bandwidth reservation




- Shared bandwidth (“flat” network)
- Security issues (access control, authentication, data encryption)
- Roaming issues (multiple Access Points and MAC address migration between ports)
- Management issues
- Antennas (coverage and beam patterns)
- Wireless repeaters and bridges



- ISDN, POTS, Leased Line
- Frame Relay
- ATM
- MPLS
- “Dark fibre” and Wave Division Multiplexing
- SONET / SDH etc.
- ADSL / SDSL
- VPNs
- Managed services (usually TCP/IP based)
- Encapsulation and tunnelling
- FC over IP, FC over Ethernet




- Routers do not need to be involved in the normal inter-node traffic within a LAN, other than keeping track of who's where and making themselves known
- Routers build knowledge of address (node or interface) reachability on a per-protocol basis
- Protocol address based (Layer 3)
- Need to design addressing scheme
- Bandwidth control
- Design routing paths
- Routing table updates are propagated between routers




Routers are generally used to interconnect LANs over a
WAN

- Separate devices or can be integrated into the core
- Need to design protocol addressing scheme and areas
- Good control over bandwidth
- Layer 3 devices – protocol address based
- IPV6 is common in big core routers
- Rare to find DECnet routing in modern routers – it's a TCP/IP dominated world in the WAN




Storage networks


- Fibrechannel technologies (1 / 2 / 4 / 8 ... Gbps)
- Storage devices (disc arrays and tape drives)
- Host Bus Adapters (HBAs) and FC switches
- WWIDs and WWNs
- SAN Segmentation (switching, routing, zoning etc.)
- Storage subsystems and device presentation
- SAN extension (eg: FC over IP)




- A switch based network optimised for shifting large quantities of data with high throughput and low latency
- All endpoints uniquely identified with a WWID (World-Wide ID)
- Multiple switches can be interconnected
- Inter-switch links can be trunked
- The network between the storage devices and the systems is known as a fabric
- Large fabrics need to be segmented
- High availability typically uses a dual-fabric SAN




- EVA family and XP family from HP
- Array controllers interface to the SAN
- Mirrored cache in the paired array controller
- Paired array controller drives the discs
- Array controllers presents devices to the fabric once Vdisks are created in the array
- Array controllers map storage to physical discs
- Array controllers managed by SAN appliance
- Systems can interact with the controller by means of scripting through the SAN appliance (eg: snapshots, snapclones, mirrorclones)




- WWIDs are unique
- Systems and storage controllers scan the fabric to build a list of paths between devices
- Storage devices (eg: EVA Vdiscs) are presented to the fabric by the array controller
- Device presentation can be controlled to limit access to specific paths (by WWID)
- Devices are presented to the fabric by the storage controller with a LUN (logical unit number) and (required by some operating systems) a device identifier (OS unit ID)



- Device paths and visibility can be controlled by zoning in the switches
- Zones can be port based, or WWID based (“soft zoning”)
- Zones can overlap (think Venn diagrams)
- Current zoning best practice uses the “single initiator, multiple targets” model
- Systems (HBAs) need to have BIOS type support for booting from SAN devices
- See the HP SAN design reference guide




- Inter-site links can be “trunked” (as with data networks) to provide sufficient bandwidth
- Link “glitches” will cause fabric resets and rescans, so use FC routing in large extended SANs to minimise disruption
- Zones can extend across multiple switches (as with VLANs)
- Wave division multiplexing (DWDM, CWDM) can be used for extended distance inter-switch links
- “FC over IP” can be used to link SANs over an IP data network (beware latency issues – use QoS techniques)




Part 2:

- Common data network protocols (TCP/IP family, NetBIOS / NetBEUI, DECnet, SCS, LAT, MOP, etc.)




Typical network protocols:


- TCP/IP (and all its components – DHCP, BOOTP, TFTP, FTP, Telnet, HTTP, SSH, NFS etc.)
- NetBIOS / NetBEUI
- DECnet-Plus (DECnet over IP) or DECnet Phase IV
- SCS (OpenVMS clustering)
- LAT (DECserver terminal access etc.)
- MOP + Remote Console (DECserver, LAVC boot etc.)
- DTSS (can be disabled)
- LAD + LAST (Infoserver etc.)




- Predominantly TCP/IP
- NetBIOS / NetBEUI and NetBIOS over TCP
- WINS naming service
- Windows file sharing
- Windows printer sharing
- Pathworks provides DECnet for Windows and Windows file & printer sharing, largely replaced by CIFS (based on Samba)
- “NIC teaming” for availability




- Predominantly TCP/IP
- “NIC teaming” for availability
- DECnet for Linux project
- SAMBA for Windows file & printer sharing




- SCS for clustering
- LAD/LAST for disc serving (InfoServer)
- Pathworks (Advanced Server) - DECnet for Microsoft Windows and LANmanager functionality for OpenVMS systems. Replaced by CIFS (based on SAMBA)
- LAT / MOP for DECservers
- “LAN failover” for improved LAN availability (all protocols). Equivalent to NIC teaming
- “failsafe IP” for improved TCP/IP availability within a cluster




- Layer 4 – port or socket layer (eg: HTTP = port 80, “well known” ports allocated by convention)
- Layer 3 – IP addressing and routing layer (eg: 192.168.0.n/24, DNS/BIND resolver user to convert IP hostnames to interface IP addresses)
- Layer 2 – MAC address layer (ARP used to convert IP interface addresses to MAC addresses, cached locally)
- Layer 1 – Physical layer (transmission media)




- DNS and the BIND resolver
- DHCP address provision
- BOOTP services
- TELNET access
- FTP / TFTP file transfer
- NFS file serving
- Monitoring with SNMP
- SMTP / POP / IMAP
- Secure extensions: SSH, SSL, IPSEC
- Printing (LPR / LPD)




- **MAC Address formed from Node address:**
 - Area 1 - 63, Node: 1 - 1023
 - 16 bit address = (Area x 1024) + Node number
 - SCSSYSTEMID = same 16 bit value
 - AA-00-04-00-nn-mm
 - nn-mm = byte reversed hexadecimal 16 bit address
- Sets MAC address on all LAN adapters based on DECnet node address, so cannot connect multiple LAN adapters to the same LAN (or extended LAN)




- The obvious big difference - NCL in place of NCP
- “DECnet over IP”
- Phase IV compatible addressing on a per NIC basis – can thus select which NICs change their MAC address
- Multiple path behaviour (multi-homed End System)
- Permanent database is NCL script files (text)
- Name Services (local, domain, DECdns)
- Time Synchronisation Service (DTSS)
- Routing algorithms (Phase V routers)



- session control
 - applications and ports
- transports
 - NSP and OSI (plus OSI templates)
- routing
- routing circuits
- csma-cd station (ethernet and FDDI)
- <datatype> links (HDLC, DDCMP etc.)
 - <datatype> link logical station
- modem connect lines




- Preserves DECnet APIs for existing applications
- Performance and availability are determined by underlying IP network infrastructure
- DECnet uses TCPIP as a pseudo-transport layer
- Need to have RFC1006 and RFC1869 (aka RFC1006-Plus) OSI transport templates - ports 102 and 399
- Streams interface
- Need to have PWIP driver enabled




Part 3:


- Network infrastructures - putting it all together




	Before	Now	After
Environment			
System			
Component			




- Bandwidth – determines throughput
- Latency – determines response time
- Jitter (“div latency” or variation of latency with time) – determines predictability of response
- Think about a motorway :
 - What is throughput?
 - What is latency?
 - What is jitter?




- Bandwidth – determines throughput
 - Worst case is the maximum capacity of the smallest path in the system (the “bottleneck”)
 - It’s not just “speed”, it’s throughput in terms of “units of stuff per second”




- Latency – determines response time
 - Worst case is the sum of all the delays in the system
 - Latency determines how much “stuff” is in transit through the system at any given instant
 - “Stuff in transit” is the data at risk if there is a failure




- Jitter (“div latency” or variation of latency with time)
 - Determines predictability of response, ie: timeouts
 - Ideally zero
 - Wildly fluctuating latency is a “bad thing”
 - Latency fluctuations can cause system failures under peak load



- Availability is more important than performance
- Scale network so that desired overall system performance is based on minimum essential number of paths and maximum estimated traffic
- May wish to take advantage of installed bandwidth capacity to provide additional functionality when everything is working
- Document and diagram the network configurations
- Install local UPS for network hardware




- Duplicate devices such as network printers and terminal servers on different paths
- Segment the network to provide gradual degradation rather than wholesale failure
- Have a fallback plan for getting to remote site equipment (eg: dial-up modem to console port of router)
- Network management is all about problem identification and rectification
- Learn the warning signs for common problems




Relative time is more useful than absolute time


- Need to be able to order events across the network based on timestamps
- UTC Timestamp format
 - Time value
 - Inaccuracy component
- External reference clocks
- NTP
- DTSS




- **Signal path quality and reliability**
 - Retransmits severely affect overall throughput
- **Bandwidth – determines throughput**
 - Large packets shift more data with less overhead
- **Latency – determines round trip delay**
 - Determines how much data is in transit at any given instant
 - Data in transit is at risk if there is a failure
- **Jitter (“div latency” or variation of latency with time) – determines predictability of round trip delay**
 - Understanding jitter is important for establishing timeout values
 - Severe latency fluctuations can cause system failures



- Traffic flow, end-to-end packet delivery, delivery failure notification and performance are key parts of the design of any network protocol, as are the addressing scheme and the naming scheme
- Multicast packets are inherently “fire and forget”
- Multiple paths through a network bring additional complications – packets may no longer arrive in the order in which they were sent
- What happens when paths fail or are intermittent?
- How do we cope with bad latency or jitter?
- Time synchronisation across the infrastructure





- Node naming, addressing schemes and routing mechanisms
- Multiple NICs and multiple LANs
- Map functions to NICs:
 - Management (ILO, SAN appliance, etc.)
 - Clustering
 - Network backups
 - Data transfers (eg: FTP, NFS etc.)
 - Interactive users




- LAN failover / NIC teaming
- TCP/IP - Failsafe IP
- DECnet Phase IV and V - load balancing
- SCS – stopping and starting per adapter with SCACP or LAVC\$START_BUS / LAVC\$STOP_BUS

- MOP and LANCP (network booting)
- LAD / LAST (InfoServer)
- LAT (DECservers)







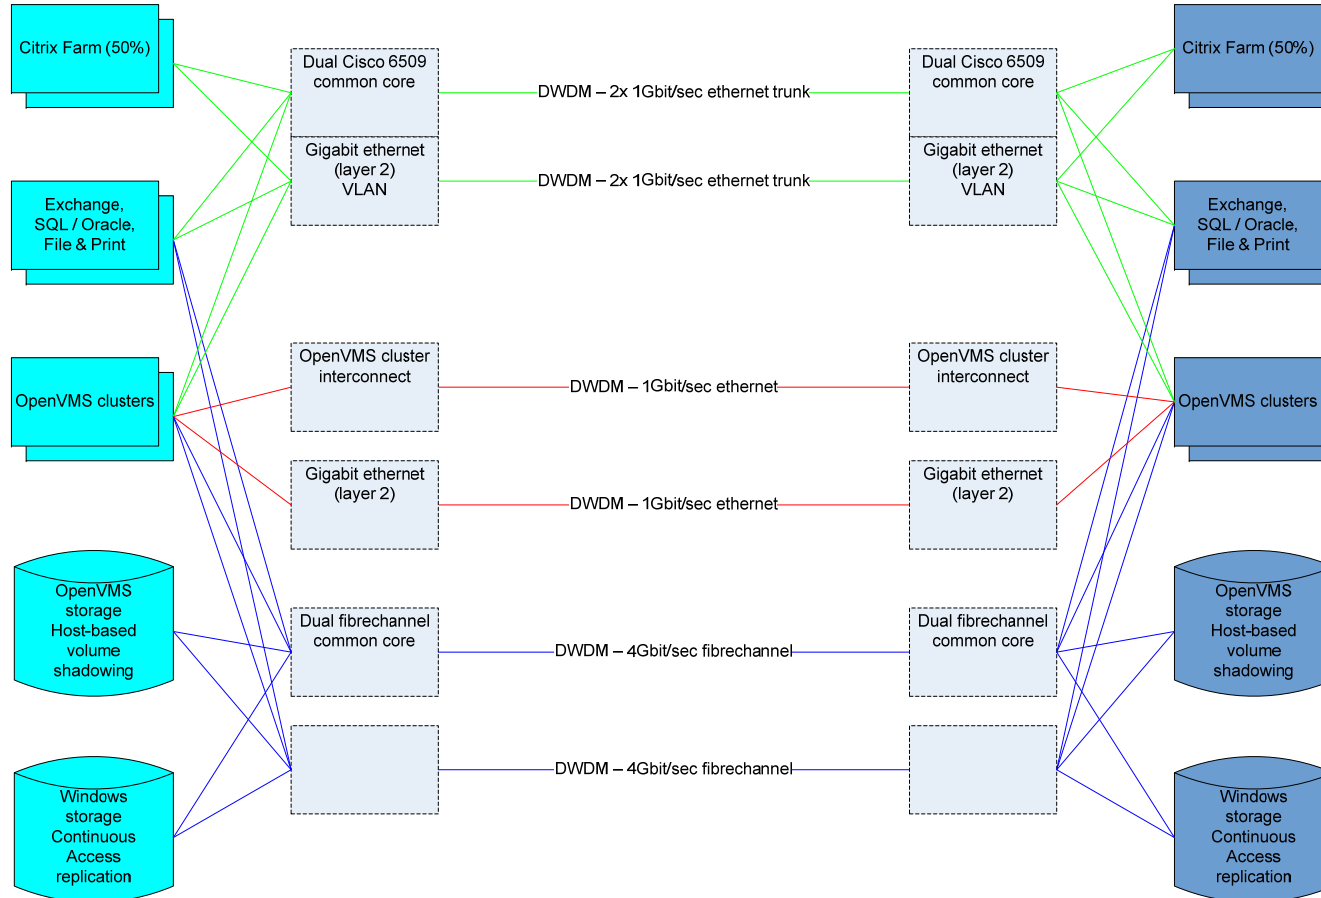
- TCP/IP – multiple NICs per subnet, dynamic routing
- DECnet Phase IV– L1 routers or end-node failover
- DECnet-Plus –Multi-homed ES or IS, load balancing



Part 4:

- **Safety-critical and mission-critical system:**
 - Migrate from Alpha to Integrity
 - Move from 3x regional clusters to single national cluster
 - Move from HSG80s to EVA4100s
 - Move to multiple NIC connectivity
- Similar principles apply in many other cases





***failsafe IP*:**



- user access
- system admin remote access
- DECnet over IP
- interface status "ping tests"



***dual rail*:**

- SCS (jumbo frames, PE driver compression and checksumming)


***LAN failover* / *NIC teaming*:**

- SSSU scripting (TCP/IP)
- DTCS monitoring (AMDS, TCP/IP)
- quorum adjustment (AMDS)
- iLO and switch access (TCP/IP)
- LAT
- DECnet







- “Converged ethernet” – fibrechannel and ethernet protocols on the same physical carrier with common interfaces and switching infrastructure
- 10GigE (and faster)
- Fibre, not copper (transmission characteristics matter)
- RNIC – offloading the bulk of the protocol handling to the NIC and minimising both CPU overhead and moving data around between the memory subsystem and the NICs




- System configuration – the art is to select components that work well together and which provide the bulk of what you need with minimal additional work
- Establish the minimum requirements that have to be met – and do it as well as possible
- Availability and performance have to be designed in
- Monitoring and automation are key components
- Understand the typical behaviour of your systems and be aware of changes
- Configuration control – hardware, firmware, settings, software etc.




- Use multiple systems within a site – and have appropriate physical separation between them
- Use multiple sites with appropriate geographical separation
- Understand what happens when a failure occurs and how the overall system configuration is likely to respond
- Avoid the risk of more than one system thinking that it's in charge when a failure has happened
- Ensure that the surrounding infrastructure and environment is appropriate to the needs of your systems
- Configure the individual components of the systems in a manner that maximises interchangeability




- Losing data is a disaster
- Availability is more important than performance
- Size storage subsystem based on minimum components and maximum estimated throughput
- Segment storage subsystem to provide gradual degradation rather than wholesale failure
- Need adequate backup capacity and throughput in order to meet permissible backup windows
- Understand application behaviour and storage performance requirements (bandwidth and latency)




- Scale network so that overall performance is based on minimum essential number of paths and maximum estimated traffic
- May wish to take advantage of installed bandwidth capacity to provide additional functionality when everything is working
- There is no such thing as a single protocol network
- Understand the behaviours of the different protocols under failure conditions
- Segment the network to provide gradual degradation rather than wholesale failure



- Think big, implement small, expand as needed
- Documentation
- Modularity
- Reliability
- Scalability
- Configuration control
- Testing (development, pre-deployment, post-deployment)
- Installation
- Monitoring
- Management and firmware updates



- Understand the basic principles by which networks work
- Network hardware tends to have a long life-cycle
- Minimise the chance of making mistakes
- Avoid complexity where possible
- Design for change over time
- Ensure that all cabling is tested and labelled
- Aim to minimise disruption when failures occur
- Use managed equipment that you can monitor easily
- Use the highest quality equipment that you can afford




SAN information:



- HP SAN design reference guide (AA-RW86N-TE)
- HP storage subsystems web site
- Brocade web site

Data networks information:

- HP ProCurve networking web site
- Cisco web site

Standards bodies





Thank you for your participation

Q & A

