

# OpenVMS TCP/IP

1E02

peter ranisch

pr - edv-dienstleistungen  
email: OpenVMS@ranisch.at

www.it-symposium2008.de

1

## Current Versions & ECO's

OpenVMS VAX V7.3	TCPIP V5.3 ECO 4	Mar '04
OpenVMS Alpha V7.3-2	TCPIP V5.4 ECO 7	Oct '07
OpenVMS Alpha V8.2 OpenVMS Integrity V8.2-1	TCPIP V5.5 ECO 3 or TCPIP V5.6 ECO 2	Feb '08
OpenVMS V8.3 (Alpha and Integrity)	TCPIP V5.6 ECO 2 (TCPIP V5.5 unsupported)	Aug '07
OpenVMS V8.3 or Later	TCPIP IPsec EAK	Aug '07

www.it-symposium2008.de

2

## Upcoming Versions & ECO's

OpenVMS V8.2 or Later	TCPIP V5.6 ECO 3*	Summer '08
OpenVMS V8.3 (Alpha and Integrity)	TCPIP V5.7†	Field Test Late '08 Ship '09

\* V5.6 ECO 3 is currently in QA

† V5.7 Under development & *requires* OpenVMS V8.3

[www.it-symposium2008.de](http://www.it-symposium2008.de)

3

## TCP/IP Version 5.6

- Shipped with OpenVMS 8.3
- OpenVMS Alpha and Integrity
- NFS server returns on Integrity
- NFS client TCP transport
- DNS/BIND 9 resolver and v9.3 server
- DNSsec
- NFS symbolic links
- NTP security update including SSL, AutoKey
- SMTP multi-domain zone
- SSH upgrade with Kerberos
- IPv6 support for printing
- FTP performance boost for VMS Plus
- Updates to TCPIP\$CONFIG (Interface menu)
- Improved management utilities (such as ifconfig)
- PPP serial-line support returns

Post V5.6 – IPsec Support (as an EAK)

[www.it-symposium2008.de](http://www.it-symposium2008.de)

4

## BIND 9 Resolver and Server

- BIND 9.3.1 for resolver and server
  - Resolver in TCPIP V5.5 was based on BIND 8
  - Server in TCPIP V5.5 was based on BIND 9.2.1
- BIND resolver
  - Lookups over IPv6
  - New ASCII configuration file (supplements existing one)
  - Improved thread support in `getaddrinfo()` and `getnameinfo()`
- BIND server
  - Includes critical updates to DNSSEC (signed zones)
  - Aligns DNSSEC with current RFCs and industry practice

## NFS Client TCP Support

- TCP transport for NFS (previously server-only)
  - Important for WAN access
  - Offers robust flow control and retransmission behavior
  - Friendly to tunneling and port forwarding

## NFS Symlink (symbolic link) Support

- A symbolic link is simply a link to another file
- When accessed, the target file is used automatically
- Deletion of the link has no effect on target file
- Links can span disks and even systems with NFS support
- Requires changes in CRTL, RMS and NFS
- NFS server must be able to create and recognize links
- NFS client must properly create, detect and follow links
- Shipped with OpenVMS V8.3
  - More updates and refinements already underway

## NTP Security Update

- Security updates from University of Delaware (UDel) NTPv4 (Version 4.2.0)
- NTP 4.2 AutoKey cryptography, using SSL
  - AutoKey is based on public key cryptography
  - Provides for secure server authentication, packet integrity, resistance against clogging and replay attacks, spoofing, and protection against masquerade.
  - Uses the OpenSSL crypto library
  - Detailed configuration steps in an Appendix of the Release Notes
  - Existing private key mechanism with MD5 remains available

## SSH Upgrade with Kerberos Support

- Kerberos support is enabled for V5.6
- DCL help for SSH commands
- SFTP/SCP
  - Improved support for additional VMS file types
    - Most popular structures are now supported
    - No support yet for RMS Indexed files
      - (You can encapsulate them in a saveset or ZIP file)

## TELNET Server Device Limit

- OpenVMS now supports large unit numbers
- Previous version (TCPIP V5.5) allowed units beyond 9999 for BG devices
- For V5.6, we added this support for TN devices

## IPv6 Support for LPD and TELNETSYM

- Allows printer communication to use IPv6
- Needed for deployment of a mostly-IPv6 network
- Note: HP enterprise printers now support IPv6

## Updated TCPIP\$CONFIG (Interface Menu)

- Previous TCPIP\$CONFIG.COM used outdated notion of cluster interfaces and one IP address per interface
- Improved configuration of multiple addresses
- Simplifies common task of changing IP address and/or hostname
- Additional information displayed to the user
- Manages both permanent database and active system
- Pseudo-interfaces continue to be stored internally

## New Look of Interface & Address Menu

```

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details:   Configured hostname=gryffindor-e0, Active=gryffindor-e0

Configuration options:

 1 - WE0 Menu (EWA0: Multimode 1000mbps) ]
 2 - 10.0.0.1/16      gryffindor-g0      Configured,Active

 3 - BE0 Menu (EBA0: Unspecified 30000mbps) ]
 4 - 1.2.3.4/8       *noname*          Configured,Active

 5 - IE0 Menu (EIA0: TwistedPair 100mbps) ]
 6 - 10.1.1.10/23    gryffindor-e0      Configured,Active

 7 - IE1 Menu (EIB0: TwistedPair 100mbps) ]
 8 - 10.1.1.11/23    gryffindor-e1      Configured,Active
 9 - 10.1.1.10/23    gryffindor-e0      Configured,Active-Standby

 I - Information about your configuration

[E] - Exit menu

```

## Interface Menu

```

HP TCP/IP Services for OpenVMS Interface WE0 Configuration Menu

Configuration options:

 1 - Add a primary address on WE0
 2 - Add an alias address on WE0
 3 - Enable DHCP client to manage address on WE0

[E] - Exit menu

Enter configuration option:

```

# Address Menu

HP TCP/IP Services for OpenVMS **Address Configuration Menu**

**WE0 10.0.0.1/16 gryffindor-g0 Configured,Active WE0**

Configuration options:

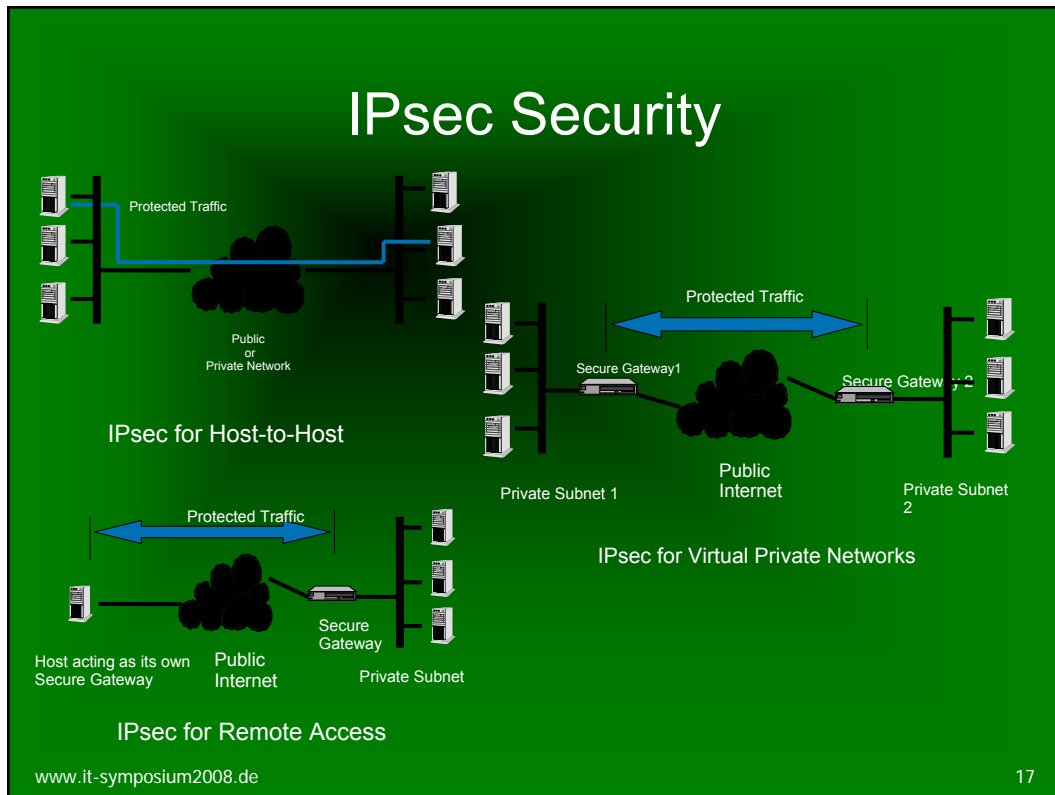
- 1 - Change address
- 2 - Set "gryffindor-e0" as the default hostname
- 3 - Delete from configuration database
- 4 - Remove from live system
- 5 - Add standby aliases to config database (for failSAFE IP)']

[E] - Exit menu

Enter configuration option:

# What is IPsec?

- Provides security at the IP layer
- Strong security that can be applied to all traffic
- Transparent to applications and end users
  - No need to train users on security mechanisms
- Protects all upper layer protocols
- Secures traffic between any two IP systems
  - Can be used end-to-end, router-to-router, or host-to-router
- Extensions to the IP protocol suite
  - Applies to IPv4 and IPv6
- Encryption and Authentication
- Key management and Security Association creation and management



## IPsec Support

- Based on the IPsec implementation from SafeNet Inc. <http://www.safenet-inc.com/> called "QuickSec"
- IPsec consists of
  - Interceptor - a platform-specific module that provides the interface between OpenVMS IP kernel and IPsec Engine module
  - Engine – a Loadable IPsec kernel module which provides crypto-processing of packets
  - Policy Manager/IKE - an application which provides processing of security policies formulated by the system manager and exchanges security policies information with remote hosts
  - Management – a set of management utilities (such as key generation, etc.)
  - Configuration tool – a basic IPsec configuration tool which processes security policies formulated by a system manager

www.it-symposium2008.de 18

## Technologies to Boost Performance

- TOE (TCP Offload Engine)
- RDMA (Remote Dynamic Memory Access)
- ETA + AIO (Embedded Transport Acceleration/Architecture) + (Asynchronous I/O)
  - PPE (Packet Processing Engines)
- Onloading (support from the chipset and the NIC manages transactions with system memory)
- PCI Express (multiple lanes, point-to-point)
- Interrupt Moderation

## Software Techniques for Performance

- Zero-copy, sendfile(), copy avoidance
- Outlining
  - Improves exception handling paths (improves branch-prediction and so avoid pipeline flushing)
- Inline with care
  - Improves deeply nested call paths
- Cloning
  - Specialised (high-performing code for common case)
- Move most common hit list items to head of list (e.g. primary addresses)
- Use hash-buckets for TCP connections
- Make network code more cache-friendly
  - I-cache
    - Reduce instruction path length
    - Improve branch-prediction to avoid pipeline flushing
  - D-cache
    - Tight data structures
    - Principle of locality (temporal proximity) improves D-cache hit-rate
- Socket Extensions and enhanced APIs
  - Provides standardised APIs to gain access to high performance features

## TOE – TCP Offload Engine

- What to offload
  - Checksumming
    - NIC performs CRC (what about errors across the bus?)
  - Large Send Offload
    - NIC performs segmentation (issues with flow control)
  - TCP/IP stack
    - NIC performs entire stack operation
      - Resource limited
      - Expensive, especially to duplicate NICs
- Existing Research shows mixed results
  - Generally unfavourable for general purpose traffic
  - Large packets provide best performance characteristics
- It's been done before (e.g. with X.25)

## Standardization

- Opengroup has defined the Extended Sockets API
  - Provides primitives for memory registration and for shared memory work queues and event queues enable applications to post operations and receive asynchronous completions **without going through the OS kernel.**

## How CPU Enhancements are Helping

- Objectives
  - Reduce OS overhead
  - Increase concurrency
  - Fewer mCPI (memory latency)
  - Shorter instruction path (better compilers, especially with improved branch prediction)
  - Multicore CPU (shorter distance between CPUs means lower latency)
- Improvement for concurrent processing of multiple network streams
  - Bigger caches
  - Shared cache between multicore CPU
    - Results in reduced memory latency as data transferred between CPUs via the caches.

## API Enhancements

- From The Opengroup (<http://www.opengroup.org>)
  - Extended Sockets API (ES-API) Issue 1.0, Jan 2005
    - Provides a shared-memory interface
    - 18 APIs
    - Improved efficiency for: buffer management, large numbers of descriptors, and supports zero-copy for transmit and receive
  - RNIC Programming Interface (RNICPI) Version 1.0, Sep 2005
  - Interconnect Transport API (IT-API) Version 2.1, Nov 2005
    - Provides an API to directly interact with RDMA capable transports
    - 68 APIs
    - Associates consumer's local memory region (LMR) and/or a remote memory region (RMR) with an Interface Adapter
- AIO (sounds a lot like QIO, with NIC polling)

## TCP/IP Services for OpenVMS Pointers and Contacts

- HP OpenVMS home page:
  - <http://www.hp.com/products/OpenVMS>
- Engineering Manager
  - [Jim.Lanciani@hp.com](mailto:Jim.Lanciani@hp.com)
- Product Manager
  - Saurabh Singhal, [Singhal@hp.com](mailto:Singhal@hp.com)
- Project Leader
  - [Mark.Hollinger@hp.com](mailto:Mark.Hollinger@hp.com)