



# OpenVMS V8.4 Update

Thilo Lauer  
OpenVMS Ambassador  
Hewlett-Packard GmbH

© 2009 Hewlett-Packard Development Company, L.P.  
The information contained herein is subject to change without notice



# Disclaimer



**„The information contained herein is subject to change without notice.**

Intel and Itanium are registered trademarks or trademarks of Intel Corporation in the U.S. and/or other countries.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.“



# Announcement OpenVMS V8.4

Date: 13 March, 2009

OpenVMS field test site,

Several weeks ago, an invitation was extended to you to participate in the upcoming field test for OpenVMS version 8.4.

HP has determined the need to continue development work on Version 8.4 and will not release a field test kit at this time. While it is difficult to give any precise dates we presently anticipate that the field test version of V8.4 will be available by the end of this year.

Thank you for your patience and understanding and continued support for OpenVMS. HP is still looking forward to your participation in the OpenVMS V8.4 field test and we hope you will continue to participate in this program.

Sincerely,

Ann McQuaid & Shobha Benakatti

# V8.4 im Überblick

- “Feature Release”
  - 6-Member Shadow Sets
  - 2 TByte Volumes
  - Clusters over IP
  - Virtualization
- Hardware Support
- Performance-Verbesserungen, Bug Fixes, ...
- Sonstiges



# 6-Member Shadow Sets

```
$ SHOW DEV DSA6
```

Device Name	Device Status	Error Count	Volume Label	Free Blocks	Trans Count	Mnt Cnt
DSA6:	Mounted	0	XMBRS	49954	1	1
\$1\$DGA1: (XYZ2)	ShadowSetMember	0	(member of DSA6:)			
\$1\$DGA2: (XYZ2)	ShadowSetMember	0	(member of DSA6:)			
\$1\$DGA3: (XYZ2)	ShadowSetMember	0	(member of DSA6:)			
\$1\$DGA4: (XYZ2)	ShadowSetMember	0	(member of DSA6:)			
\$1\$DGA5: (XYZ2)	ShadowSetMember	0	(member of DSA6:)			
\$1\$DGA6: (XYZ2)	ShadowSetMember	0	(member of DSA6:)			



# 6-Member Shadow Sets

- Benutzung durch Angabe von mehr als 3 Members  
\$ MOUNT/SYSTEM DSA16 -  
\_ \$ /SHADOW=(\$1\$dga1, \$1\$dga2, \$1\$dga3, \$1\$dga4) label
- Keine neuen Qualifier
- Keine Änderungen in DISMOUNT



# 6-Member Shadow Sets

- Mixed version Support
- 3-member Shadow Sets weiterhin supported
- Benutzung von >3 Members erfordert neue Software
- ohne neue Software wird MOUNT/INCLUDE evtl. nicht alle Member finden

# 6-Member Shadow Sets

- Performance Tests laufen
  - Reads **können** schneller sein
  - Writes **werden** langsamer sein
- `$ SET SHADOW/COPY_SOURCE`





## 2 TByte Volumes

- bisher: maximale Größe eines Volumes = 1 Terabyte
- Problem:  
[www.hardwareschotte.de](http://www.hardwareschotte.de): 1TB Disk 96,47 €

# 2 TByte Volumes

## Restriktion:

- LBNs in Longword gespeichert
- maximal  $2^{32}$  Blocks adressierbar (2 TByte)
- ... oder doch nur  $2^{31}$  Blocks? (1 TByte)

## Lösung:

- Benutzung von Bit 31 bei LBN-Adressierung

# 2 TByte Volumes



**NEIN!**

# Clustering über TCP/IP

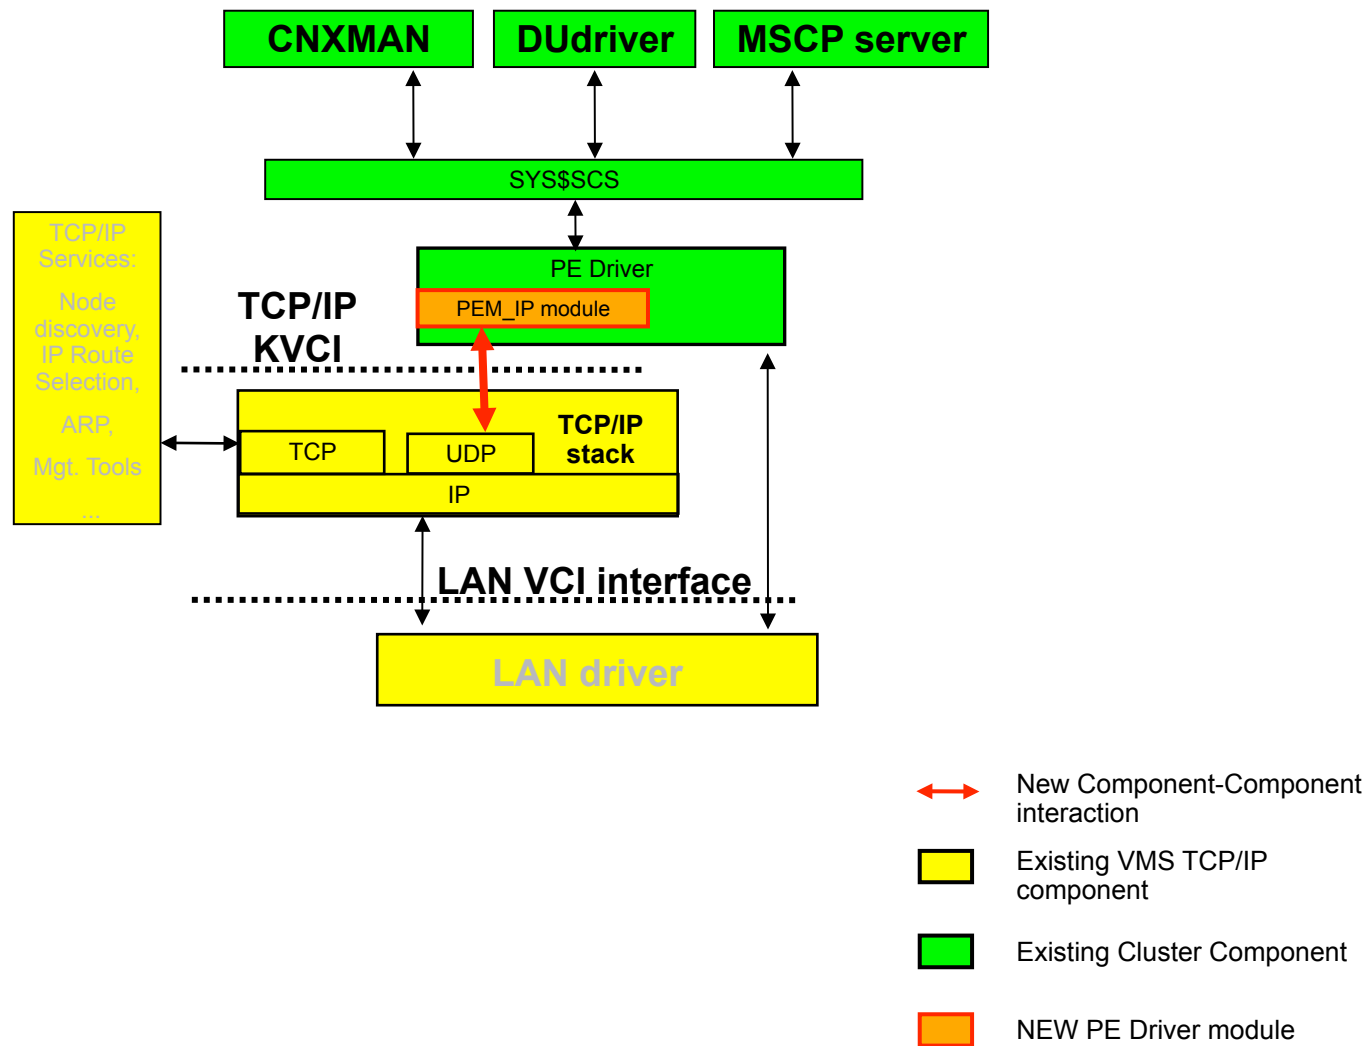
- Warum

- VMS Cluster-Protokoll ist bridged, nicht “routable”, benötigt spezielle Netzwerk-Komponenten
- Zusätzlicher Aufwand für Netzwerk-Admins
- Switch-Hersteller vernachlässigen “bridged” Support
- vereinfacht DT-Konfiguration – IP intern und extern

- Wie

- Cluster-Kommunikation wird in IP-Pakete verpackt

# IPCI-Lösung: PE Driver nutzt UDP






# VMS und Virtualisierung

- Guest System in HP VM v4
- Support für VSE Suite v4.1
- Beta Test erfolgreich verlaufen
- Neue/optimierte I/O Driver für Guest Systeme
- Agents für OpenVMS
  - Capacity Advisor, Virtualization Manager, gWLM

# Hardware Support



- “Tukwila”
  - 4 Cores
  - 2 GHz
  - 30 MByte Cache
- MSA2000fc G3
- MSL G3 Libraries
- VLS Support



# Hardware Support

- “Tukwila”
- MSA2000fc
  - 2 \* 4Gb FC Controller
  - 1GByte Cache
  - 4 LFF Drive Enclosures und 60 Drives, oder  
3 SFF Drive Enclosures und 99 Drives
  - 29.7 TByte SAS, 60 TByte SATA, 512 LUNs
- MSL G3 Libraries
- VLS Support



# Hardware Support



- “Tukwila”
- MSA2000fc G3
- MSL G3 Libraries
  - bis zu 4 LTO4 Drives
  - 24/48/96 Cartridges
  - 38/76/153 TB Kapazität
  - 1,15/2,3/3,4 TB Transferrate
  - 4 Gb Fibre Channel
  - LTO1/2/3/4, SDLT<sub>x</sub>
- VLS Support



# Performance



- Exception/Unwind
- “Lazy” FP Save/Restore
- RAD/NUMA Support
- iCache Flush
- Process VA Deletion
- XFC deferred writeback

# Unix Portability



- RMS
  - Posix Pathnames
  - symlinks
- Semaphores

# RMS – Posix Pathnames

- Support von Logical Names
- Erstes Feld in absolutem Pfad wird geprüft auf:
  - Logical Name
  - Directory oder File Name in Posix Root
  - VMS Device
- Beispiele:
  - `/SYS$LOGIN/login.com`
  - `/DKA0/test/file.txt`
  - `/bin/GCC.EXE`

# RMS – Symbolic links

- Ein symbolic link ist ein Directory-Eintrag, der einen Namen mit einem Text-String verbindet
- Der Text-String wird von RMS als POSIX Pathname interpretiert
- in OpenVMS implementiert als File mit Organization SPECIAL und Typ SYMBOLIC\_LINK
- Symbolic link mit relativem Pfadnamen wird relativ zum enthaltenden Directory verfolgt
- Absolute Pfadnamen in Symlinks: ab root, Mount Point, Device



# RMS – Symbolic links

```
$ CREATE/SYMLINK="a/HP_buys_IBM.txt" IBM_buys_HP.txt
```

```
$ CREATE [.a] HP_buys_IBM.txt
```

Rumours, just Rumours...

```
^Z
```

```
$ DIR/DATE IBM_buys_HP.txt
```

```
IBM_buys_HP.TXT -> a/HP_buys_IBM.txt
```

# RMS – Symbolic links

- Voller Support in System Services
  - `$OPEN`, `$CREATE`, `$SEARCH`, `$ERASE`, `$RENAME` ...
- Voller Support durch C RTL
  - `symlink`, `readlink`, `unlink`, `realpath`, `lchown`, `lstat`
- Neue Lexical Functions verfügbar
  - `F$READLINK` `F$SYMLINK_ATTRIBUTES`
- Symlink Support für Commands, die Filenames akzeptieren
  - Meist wird per Default dem symbolic link gefolgt
  - Ausnahmen: `BACKUP`, `DELETE`, `DIRECTORY`, `PURGE`, `RENAME`
  - `/[NO]SYMLINK` Option für `COPY`, `DIRECTORY`, `DUMP`, `SET FILE`
  - `/EXCLUDE` wirkt direkt auf den symbolic link

# Semaphoren

- Zugriff über C RTL
- Posix Semaphores
  - `sem_open()`, `sem_post()`, `sem_wait()`, etc.
- System V Semaphores
  - `semop()`, `semctl()`, `semget()`
- Thread-aware



# TCP/IP V5.7



- Packet Processing Engine (PPE)
- weitere Verbesserungen
- IPSEC?

# Performance, Performance, Performance



- 10 GB Ethernet ist Realität
- Probleme mit High-Bandwidth Feeds
  - Protokoll-Overhead
  - Memory Latency
    - Viele Buffer-Copy-Operationen zw. NIC und App.
  - Interrupt Load
  - Instruction Path Length
  - Multiple-passes over data
  - CPU-Auslastung
  - Cache Misses
  - CPU kann saturieren durch Packet Processing

# TCP/IP PPE: Konzept

- Vorbild: OpenVMS Dedicated Lock Manager
- Bisher:
  - TCP/IP läuft auf timesharing CPU
  - bei saturierten CPUs kann TCP/IP zum Bottleneck werden
- PPE bindet CPU zu TCP/IP
  - kein Sharing mit Prozessen
- PPE dynamisch ein-/ausschaltbar
  - Kundenspezifisches Monitoring, PPE scriptgesteuert

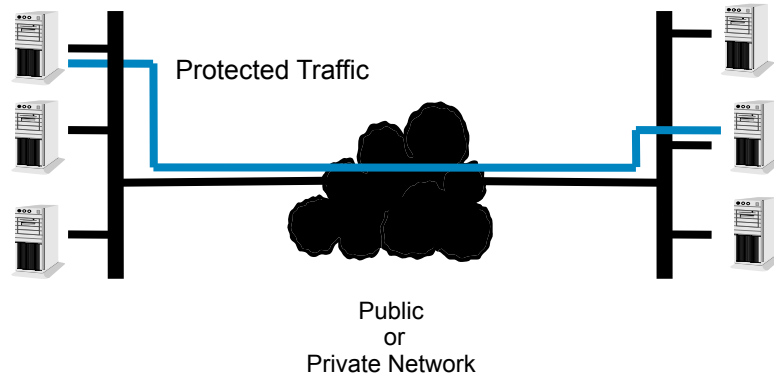
# TCP/IP PPE: Vorteile

- Keine Context Switches
- Keine Cache-Konflikte mit Applikations-Prozessen
- **Polling** vermeidet Interrupt Overhead
- Kein Fork Processing

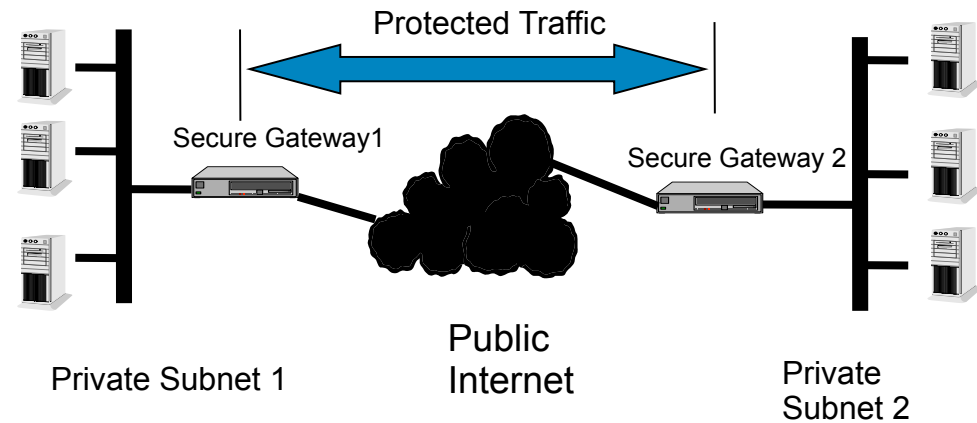
# What is IPsec?

- 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

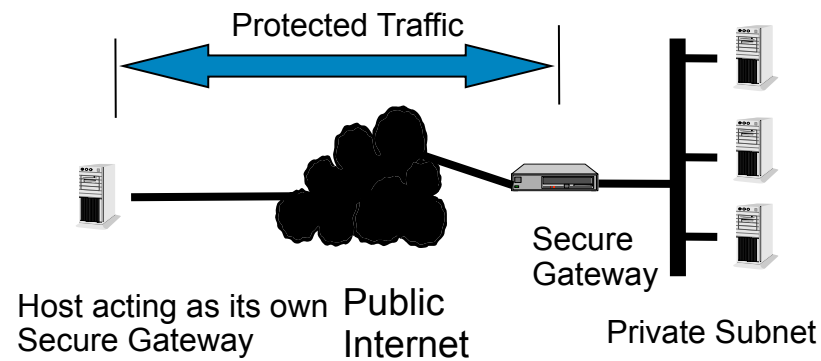
# IPsec Security



IPsec for Host-to-Host



IPsec for Virtual Private Networks



IPsec for Remote Access

# IPsec Status

- EAK available now
  - <http://h71000.www7.hp.com/openvms/products/ipsec/index.html>
- EAK includes PPE
- Firewall function included in IPsec code base
  - No separate firewall product planned

# System Management

- weitere Provider
  - Storage Controller/Media (Events, Thresholds)
  - Netzwerk (DNS, LAN, Ports)
  - Insight Power Manager (Verbrauch, Regelung)
  - Blade HW Support
- Provisioning-Erweiterungen
  - Lizenz-Management via HP SIM
  - Cluster & TCP/IP Konfiguration via HP SIM
- Infrastruktur
  - OpenPegasus V2.7
  - GiCAP
  - OpenView-Erweiterungen





# Sonstiges



- Security
  - OpenVMS nun digital signiert
  - LDAP External authentication
  - SSL Refresh (OpenSSL 0.9.8H + Security Fixes)
  - Stunnel 4.2 verfügbar
  
- CIFS V1.1 ECO 1
  
- ABS V4.5
  - Hardware + Software Encryption
  - Key Management

# OpenVMS Cluster Support

## Matrix



Integrity Server

		AlphaServer			
		v7.3-2	v8.2	v8.3	v8.4
v8.2	v8.2	Warranted	Warranted	Migration	Migration
	v8.2-1	Warranted	Warranted	Migration	Migration
	v8.3 & v8.3-1H1	Migration	Migration	Warranted	Migration
	v8.4	Migration	Migration	Migration	Warranted

**Notes:**

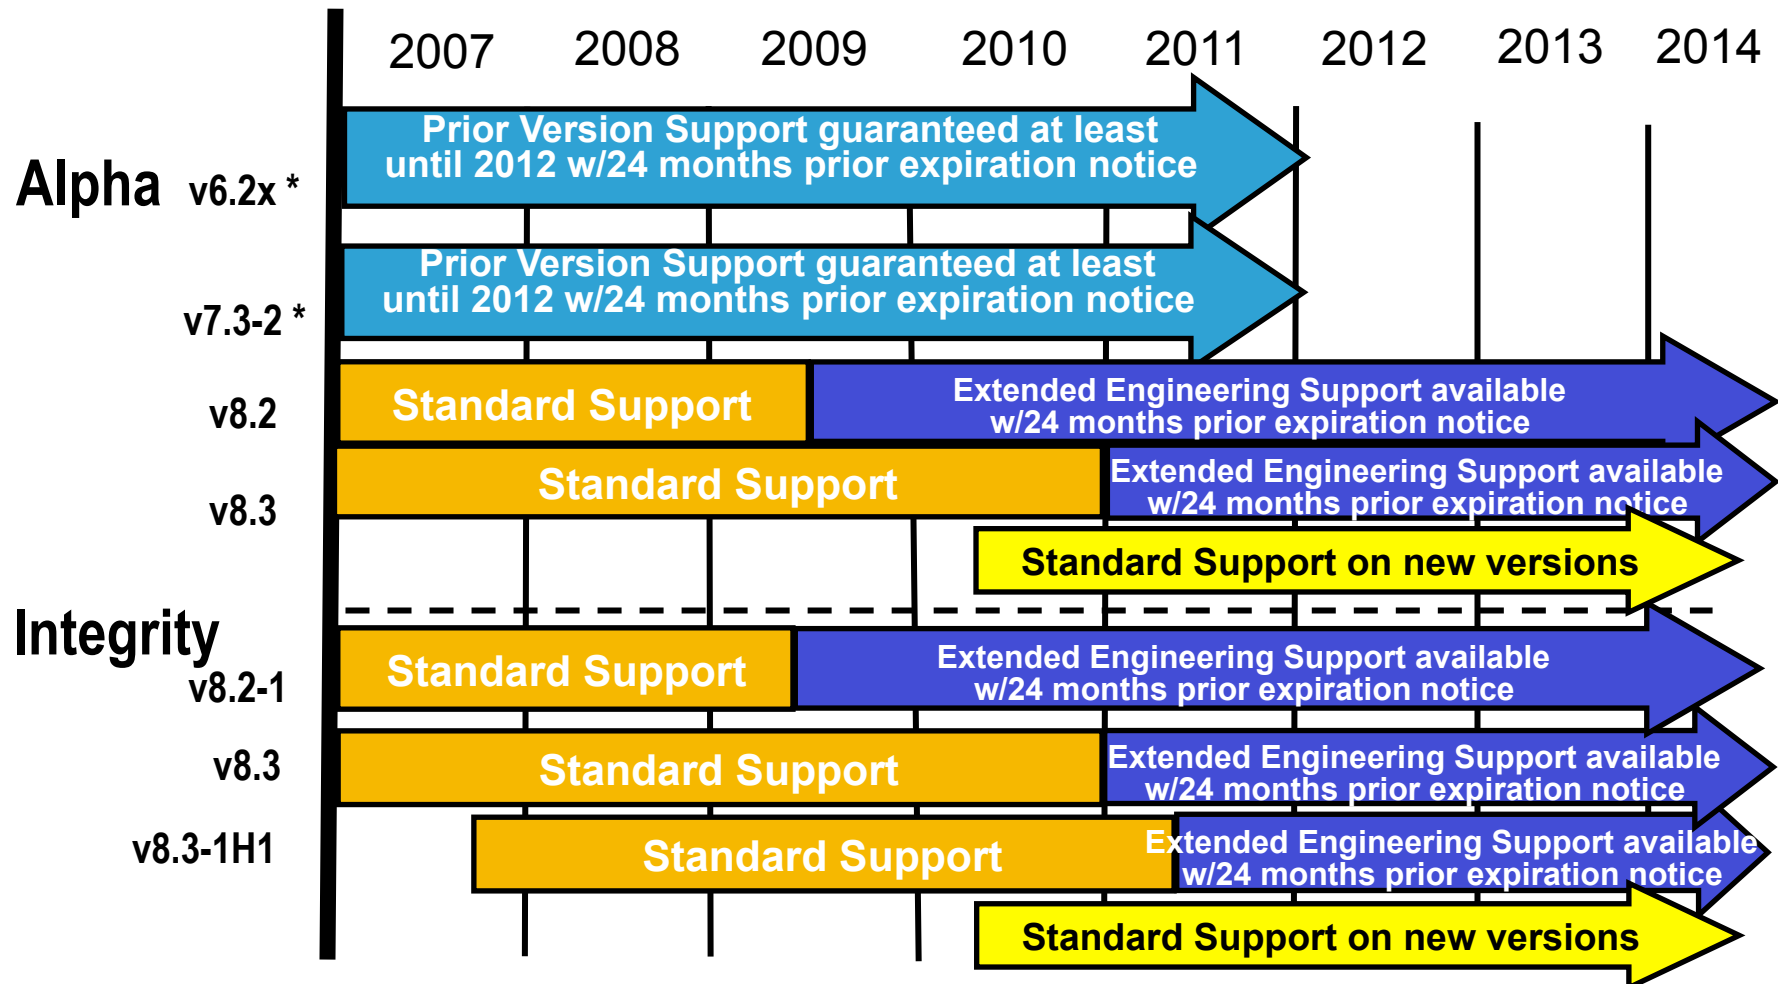
- Same platform/version pairings always warranted
- Alpha v7.3-2 and v8.2 are Warranted Together

VAX

		AlphaServer			
		v7.3-2	v8.2	v8.3	v8.4
v7.3		Warranted	Warranted	Warranted	Warranted

**Note: VAX and Integrity Servers can exist in the same cluster ONLY for temporary migration purposes**

# OpenVMS Service Support Roadmap



• Prior Version or Standard support will be provided on these versions at least until 2012.

\*\* w/24 mo notice: A 24-month notification will be provided before support is ended.

*Standard support ends when the 2nd subsequent release ships. HP supports the current version and one back.*

*Extended Engineering Support will be available for additional cost. Future version shipment dates are estimates.*

# Kontakte



- Gaitan D'Antoni (OpenVMS Technical Director)
  - [Gaitan.Dantoni@hp.com](mailto:Gaitan.Dantoni@hp.com)
- Leo Demers (OpenVMS Product Management)
  - [Leo.Demers@hp.com](mailto:Leo.Demers@hp.com)

