

MIF-arch-dt status update: the Multiple Provisioning Domain (MPVD) Architecture

Dmitry Anipko (architecture document editor)

Tim Chown (presenter, on behalf of the design team)

IETF 88 MIF WG

Vancouver, 7th November 2013

MIF-arch-dt activities...

- Multiple Provisioning Domain (MPVD) Architecture draft updated **draft-anipko-mif-mpvd-arch-05** published at the IETF opening
 - This is the focus of this talk
- Charter update proposal sent to the MIF WG list on 28th October
 - Separate charter discussion to follow after we discuss the Architecture text
- Strawman drafts for adding PVD support via DHCP and RA/RS published as **draft-kkb-mpvd-dhcp-support-00** and **draft-kk-mpvd-ndp-support-00** on 23rd October
 - Again, separate presentations on these to follow later in this MIF session

The problem

As documented in RFC 6418 (MIF problem statement)

1. Lack of consistent and distinctive management of configuration elements, associated with different networks.
2. Inappropriate mixed use of configuration elements, associated with different networks, in the course of a particular network activity / connection.
3. Use of a particular network, not consistent with the intent of the scenario / involved parties, leading to connectivity failure and / or other undesired consequences.

Goals of the architecture text

The architecture document proposes a solution by

1. Introducing a formal notion of the PVD, including PVD identity, and ways for nodes to learn the intended associations among acquired network configuration information elements.
2. Introducing a reference model for a PVD-aware node, preventing inadvertent mixed use of the configuration information, which may belong to different PVDs.
3. Providing recommendations on PVD selection based on PVD identity and connectivity tests for common scenarios.

MPVD Architecture status recap from IETF 87

- A PVD is a consistent set of network configuration information
- A PVD-aware node can associate network configuration information with one or more PVD(s) and use the PVD(s) to aid connectivity decisions
- A node may learn PVD information explicitly or implicitly
 - Explicit PVDs require a mechanism to pass information to PVD-aware clients, and may span multiple interfaces
 - Implicit PVDs are formed from information received from multiple interfaces where a PVD aware node receives no PVD information from the network
- A PVD consists of PVD elements, e.g.,
 - IP address, DNS resolver, HTTP proxy server, etc.
- A node may learn about multiple provisioning domains

Architecture status recap ctd...

- Classically, all info available on one interface is provided by a single source (e.g., net admin) and can be treated as one PVD
- Multihoming can result in more than one PVD on a single link
- An explicit PVD can span multiple interfaces and links
- Each PVD has an ID
 - Explicit PVD has globally unique ID passed to it, possibly human readable
 - Implicit PVD generates a globally unique ID
 - Choice of how IDs are advertised is left to the provider
- Dual-stack PVDs supported; PVD may have multiple address families
- PVDs may be trusted or untrusted
- PVD-aware applications will have an API for PVD selection

MPVD arch draft updates since IETF 87

- As per discussion in IETF 87
 - No hierarchical relationship assumed between interfaces and PVDs
 - Hosts can augment PVDs with information, if explicitly configured to do so; but default auto merge of different PVDs is not safe and is not recommended (see also RFC 6418)
 - Connectivity tests recommended before/during use of PVD
 - Relationship to connection managers described (see also RFC 6419)
- Section 3 added for considerations for support in DHCP, RA/RS, IKEv2
- Security considerations section added

Any other questions?

Please give feedback on the MPVD architecture text

Is it heading in the right direction?

Is there anything missing?

Any bad assumptions?

Mail list is mif@ietf.org, see <https://www.ietf.org/mailman/listinfo/mif>