

# Towards the Secure Initialisation of a Personal Distributed Environment

Scarlet Schwiderski-Grosche, Allan Tomlinson, David B. Pearce

Technical Report  
RHUL-MA-2005-9  
20 July 2005



Department of Mathematics  
Royal Holloway, University of London  
Egham, Surrey TW20 0EX, England  
<http://www.rhul.ac.uk/mathematics/techreports>

## Abstract

The Personal Distributed Environment, or PDE, represents a new concept of a computer network that takes a user-centric view of personal devices to create a purely virtual personal network. A PDE integrates all devices that are relevant to a user, regardless of their location, current role, or their specific capabilities. In this paper we present an overview of the PDE concept and describe how such an environment may be constructed. Since the user's devices may be located in several different security domains, our focus is on the secure initialisation of the PDE and to that end, we introduce a step-by-step procedure describing how user's devices will join the network in a secure manner.

## 1 Introduction

A traditional computer network may be thought of as a collection of computers, physically connected by wires, so that they can communicate with each other and share resources. The fact that all devices in such a network are connected by wires, offers an inherent level of security: these devices can only communicate if they are physically connected. Shielding physical access to a computer network, or perimeter security, offers protection from intruders. Moreover, traditional computer networks employ security gateways, such as firewalls and intrusion detection systems, at the edge of the network to implement security measures and to monitor all traffic that accesses the network.

To provide remote access to their internal network, many organisations employ Virtual Private Network (VPN) technology. VPNs send data over the public Internet through secure tunnels. In a VPN, a level of security corresponding to traditional networks is achieved through encrypting and encapsulating the transmitted data, for example using IPsec [5].

More recently, wireless communication technologies, such as IEEE 802.11 [13] or Bluetooth [3], are being used to extend the traditional networking concept even further with mobile devices that join and leave a computer network spontaneously. The security of wireless access points is problematic because the wireless nature subverts the physical security mechanisms of a traditional network [4].

VPNs and wireless communications represent two reasons why the concept of the computer network is becoming less well defined, and why new techniques have to be developed to achieve a level of security comparable to traditional computer networks.

With the Personal Distributed Environment (PDE) [6, 19], we go one step further and create a dynamic and *heterogeneous* network that exists purely on a virtual level. It is constructed from all the devices that relate to one specific user. That is, devices that are owned by the user, to which they have access rights, or through which they receive data or services. The PDE is, therefore, a *user-centric* implementation of a global dynamic and heterogeneous communication network.

The heterogeneous nature of the PDE arises from the fact that consumers will purchase devices to meet a wide range of needs. Such consumer devices are normally used with limited connectivity in the different living environments of the user. For example, their home, office, car, or personal surroundings. These devices may be fixed or mobile, and may communicate using a broad range of wired or wireless communication technologies. Furthermore, the devices' capabilities and their role in the PDE differ tremendously.

Given this environment, the task of constructing the PDE and managing connectivity presents a significant research challenge. Providing security within the PDE presents a further challenge for the following reasons:

- The PDE is a purely virtual network. Physical security mechanisms and dedicated security gateway technology are, therefore, not applicable. This means that security mechanisms need to be implemented purely in software and, furthermore, need to be ubiquitous throughout the PDE.
- A PDE consists of devices originating in different environments which implement different legacy security systems. In practice, then, any PDE security mechanisms must take existing security systems into account. The PDE cannot, therefore, impose any security mechanisms, but must support a wide range of legacy systems.
- The PDE incorporates wireless as well as wired communication technologies. Wireless communication is inherently insecure and the PDE incorporates a wide spectrum of wireless technologies which results in a heterogeneous environment.
- The PDE is dynamic in nature, the user and many of their devices are mobile. Consequently, not all devices will be available at all times.

To create a PDE the construction of a mapping table is proposed, that maps physical devices in different environments onto logical PDE components. This mapping table primarily contains information on devices, and management information for the PDE. The information in the mapping table

must be made available to all parts of the PDE. Once this basic information is available, the relevant security parameters, that fulfill the desired security services within the PDE, can be determined. With this paper, we take a first step towards designing a security architecture for the PDE by considering how the mapping table is created, and how it is filled with relevant information when the PDE is initialised.

The paper is structured as follows: Section 2 summarises related work, and section 3 then presents an overview of the PDE concept. Since the concept of the PDE and its objectives are new, and differ in many respects from existing applications of virtual networks, section 4 formulates a number of general assumptions about the PDE. We use these assumptions as a basis for the initialisation procedure described in section 5. In addition to adding of new devices to create the PDE, devices also have to be removed, for example, if they are compromised or become obsolete. This is briefly addressed in section 6. Section 7 contains conclusions and an outlook on future work.

## 2 Related Work

A number of research initiatives exist that are similar to the work on the Personal Distributed Environment being carried out by the Mobile VCE<sup>1</sup>. For example Grid Computing, Pervasive Computing, and more recently, the IST MAGNET project<sup>2</sup>. Since the PDE is a purely virtual network, we looked at these other research initiatives focussing on virtual networks, mainly Grid computing systems which are built on the concept of the Virtual Organisation. What the PDE offers in addition to Grid and Pervasive computing, is that the PDE is primarily a user centric, personal network, as opposed to a multi-user computing resource. Furthermore, the PDE is focussed on communications between personal mobile devices rather than the interaction with smart spaces envisaged by Pervasive computing, or the sharing computational resources and data, as is the objective of the Grid community. The objectives of both the MAGNET project and the Mobile VCE work follow the vision expressed by the Wireless World Research Forum<sup>3</sup> (WWRF). The WWRF is a global collaboration of industrialists and academics who envisage a concept like the PDE, known as the MultiSphere [14]. The scope of the IST MAGNET project is closest to that of the Mobile VCE, but this is relatively new work, the project being launched 2004.

---

<sup>1</sup>[www.mobilevce.com](http://www.mobilevce.com)

<sup>2</sup>[www.telecom.ece.ntua.gr/magnet/](http://www.telecom.ece.ntua.gr/magnet/)

<sup>3</sup>[www.wireless-world-research.org](http://www.wireless-world-research.org)

**Grid Computing:** The concept of a personal virtual network has strong parallels with that of the Virtual Organisation, the model used in computational Grids [7, 8]. This technology represents a mature, albeit wired, virtual environment. From a security perspective, however, many of the problems arising in Grid technology also arise in the PDE. For example, both environments are faced with the challenge of integrating numerous isolated security domains where each domain may have its own security policy and run specific security protocols. A comparison of the PDE and Grid environments is given in [20], which suggests how Grid security may be applied to resolve some of the problems in the PDE.

**Pervasive Computing:** Further parallels may be found in the field of Pervasive Computing [9, 18] which, like the PDE, is concerned with delivering personalised services to users. Unlike the PDE, however, the means by which this is accomplished is through the use of smart spaces. The focus of the PDE, on the other hand, is the construction of a personal network for an individual user. Consequently, much of the PDE work concerns the details of how to construct this heterogeneous network based on the user's local Bluetooth, WiFi, or UMTS networks. This construction therefore requires some degree of interaction between the user and the various network operators [16]. However, as with Grid computing, the security issues facing Pervasive computing share many similarities with PDE.

**WWRF:** The Wireless World Research Forum has its roots in an earlier EU-IST programme, the Wireless Strategic Initiative (WSI). The WSI invited a Think Tank of industrial and academic experts to contribute to a "Book of Visions 2000" [14]. This report describes mobile communications technologies, and business models, expected to become operational in the next decade. The WSI subsequently decided to develop this Think Tank into an open forum and as a result the WWRF was launched in 2001 [15]. Working to similar timeframes and with similar objectives, the WWRF continues the work of the WSI Think Tank and has produced a revised "Book of Visions 2001" [21]. In common with the PDE, the WWRF MultiSphere is based on a user-centric approach. The MultiSphere models layers of wireless communications technologies as concentric spheres with the user at the centre. Layer 1 is the PAN environment; layer 2 includes devices in the user's immediate surroundings and subsequent layers include longer range communications. The outer layer of the MultiSphere, layer 6, visualises a "personal

cyberworld” of communications and services.

**MAGNET:** In contrast to the mature technologies of Grid and Pervasive computing, MAGNET is a new initiative. MAGNET (My personal Adaptive Global NET) is an integrated project supported within the Sixth Framework Programme of the EU Commission and started out in January 2004. The MAGNET concept of the Personal Network incorporates the same basic ideas of the Mobile VCE PDE. According to the goals described on the IST-MAGNET web site, “the MAGNET overall objectives are to design, develop, demonstrate and validate the concept of a flexible Personal Network that supports resource-efficient, robust, ubiquitous service provisioning in a secure, heterogeneous networking environment for nomadic users. Of paramount importance is the requirement that a Personal Network will support the user in both private and business activities, while safeguarding the security and privacy of the users and their data.”

### 3 Overview of the PDE

Dunlop et al. [6] describe the PDE as “encompassing a user perspective of multiple devices (both local and remote) accessing multiple services via multiple networks, all of which can be changing dynamically”. This concept goes further than just accessing content, but harnesses the PDE capabilities in novel ways to create a new and diverse range of services for the user.

PDE devices are typically those devices that play a role in the user’s private and professional environments. It is anticipated that there will be around 20 devices in a PDE, originating in different environments. Each environment forms its own PDE sub-network, for example in the

- home environment: the home PC, printer, set-top box (STB), HiFi, or even refrigerator;
- office environment: the office PC, printer, or projector;
- car environment: car radio , navigation system, and alarm; and the
- personal surroundings: PDA, smart phone, blood pressure monitor.

In describing the PDE, we distinguish between two classes of devices, home devices and foreign devices, depending on their origin:

**Home devices:** These are devices that are, at least partially, controlled by the user. There are two possibilities where a device may be considered a home device, namely:

1. The device is owned by the user. In this case, the user has administrator rights on the device. The device may be shared with other users.
2. The device is not owned by the user, but belongs to somebody else (for example, an office PC). However, the user has access rights, an account, on the device.

For home devices, although the user may be the same for a number of devices, the user ID may differ from device to device.

**Foreign devices:** These are devices that are not controlled by the user. In this case, the device is owned by a third party and the user has no dedicated account on the device. Examples are devices that deliver a service, such as display or sensor functionality, that the user wants to make available to their PDE.

Both home and foreign devices may be static or mobile, and they may utilise different access technologies, for example fixed line, cellular, WLAN, or broadcast. Some devices will support more than one access technology. PDE devices will also have a wide range of capabilities. These will vary from powerful PCs and laptops that are always on; through devices with small displays and keypads with limited processing and battery power; to personal sensor devices. Moreover, not all devices will play an active role in the PDE, for example, peripherals and sensors that simply deliver a service to the user.

PDE devices may also access remote resources across wide area networks. Hence, the PDE may be regarded as a combination of *local* devices, connected using PAN technologies such as Bluetooth<sup>4</sup> and ZigBee<sup>5</sup>; and *remote* devices connected via a range of wired or wireless communication technologies such as fixed line, cellular, WLAN, or broadcast. It should be emphasised that devices may join and leave the PDE in an *unpredictable* way, especially those wireless devices, which may be switched off or out of reach. This results in the dynamic and heterogeneous environment that distinguishes the PDE from more mature virtual networks.

Figure 1 sketches a PDE, with personal devices originating in different environments and communicating using a wide range of wired and wireless communication technologies.

---

<sup>4</sup>[www.bluetooth.com](http://www.bluetooth.com)

<sup>5</sup>[www.zigbee.org](http://www.zigbee.org)

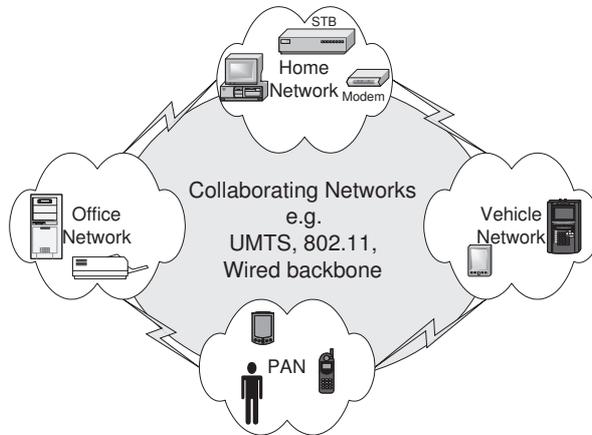


Figure 1: Overview of the PDE

The PDE provides ubiquitous access to services using all currently available personal devices and service networks. In doing so, the user preferences, and the current context, need to be considered. Accessing a service may involve accessing content on another PDE device. For example, a remote user may wish to access content such as a programme stored on a set-top box at home. Alternatively the content may originate outside the PDE, such as incoming email, voice, or broadcast video. In the delivery of services initiated outside the PDE, user preferences and context must be considered. For example, if a user's preferences state that email should be forwarded immediately when the user is traveling in a car, this could require conversion of the textual email to audio for playback through the car's onboard communicator.

Such service realisation requires knowledge of the user's preferences, currently accessible devices, and their capabilities [6]. To facilitate this, the PDE requires a **Device Management Entity**, or DME. The DME is proposed as the contact point for all incoming services, and as the location database for PDE devices not in the immediate vicinity. Consequently the DME needs to be permanently contactable. The DME intelligently directs incoming sessions using the information it possesses. The DME, therefore, maintains information on all devices in the PDE and their status.

In conclusion, the PDE represents a dynamic and heterogeneous virtual network that maintains a user-centric view of devices and services. The security of the underlying network infrastructure however, presents a number of challenges that must be overcome before the PDE can be realised. Many of the issues are described in earlier work [19]. In this paper we now propose a foundation for a secure personal network infrastructure upon which the PDE

can be based. We introduce a step-by-step initialisation procedure which goes some way towards the development of the security architecture for the PDE.

## 4 Assumptions

As the concept of the PDE is new and the architecture is still being developed, we have to make certain assumptions regarding its underlying infrastructure and functionality. In the following, we state the assumptions that have an impact on the PDE initialisation procedure and the PDE security architecture:

1. As described in section 3, PDE devices exist in a number of different environments. Each of these environment is expected to implement specific local security mechanisms. We assume that the security mechanisms within these security domains are unlikely to be compatible with each other or with any PDE specific security mechanisms. However, the local security mechanisms in sub-networks must remain unaffected by the PDE functionality. Therefore, the PDE security architecture must inter-operate with legacy security solutions.
2. We assume that PDE devices, or whole PDE sub-networks, may be unavailable at times. This applies especially to mobile devices, which may be switched off, conserving power, or simply out of reach.
3. PDE devices have varying capabilities, some of which may be very limited in nature. Hence, we assume that not all devices will be able to implement sophisticated security mechanisms, such as those based on public key cryptography.
4. We assume that some basic PDE management functionality will be available at all times. Although we do not preclude the option of the user providing this functionality themselves, we expect that this service will be provided by a server, or portal, maintained by a PDE service provider. Placing PDE management functionality on such a central server has a number of advantages:
  - The user can gain access to their PDE without the need to have physical access to any of their own devices. This may be accomplished through the Internet from any IP-enabled device such as a public terminal. Alternatively the user could access their PDE via telephone, similar to telephone banking.

- If the PDE is hosted at a service provider who can provide and maintain PDE software, the user's technical involvement can be kept to a minimum.
  - The PDE service provider may be able to offer their own services to the user, or broker PDE services from other service providers.
5. The last assumption is that the PDE should be robust. Thus, even if substantial parts of the PDE are unavailable, we assume that the remaining PDE devices and sub-networks remain operational.

The last two assumptions have further implications on the operation of the PDE.

While these assumptions are essential to the PDE architecture they are, to some extent, conflicting. Assumption 4 states that some PDE functionality should be available on a central server. On the other hand, assumption 5 states that the PDE should be robust. Since the PDE is organised in a number of sub-networks, this implies that sub-networks are able to operate independently if global connectivity is lost. However, if this is the case and some sub-networks have lost connectivity to the PDE server, then the PDE server cannot give a complete picture of what is going on in the rest of the PDE. As a consequence, the PDE server will have unsynchronised or incomplete data about the different sub-networks at different points in time.

## 5 Initialisation of the PDE

Having described the PDE model, and the assumptions that we make about this model, we may now describe the methods used to construct the PDE. Initialisation is a two stage process. The first stage is to create the framework to support the PDE, the second stage is to register devices to construct the PDE. Central to this discussion is the notion of the Device Management Entity, and this is described in more detail before discussing the details of the initialisation procedure.

### 5.1 The Device Management Entity

The PDE is a virtual network, and corresponding physical devices have to be mapped into this virtual network. Consequently, a mapping table must be created that contains information about the user's devices and how they are mapped into the PDE. This mapping table realises the virtual network, hence it must be globally available to the entire PDE. The mapping table

must be available to all PDE devices in all PDE sub-networks. Therefore, this mapping table is stored in the PDE's **Device Management Entity (DME)** [1].

In section 4, we stated that some basic PDE management functionality should be available on a central server. This management functionality is the job of the DME. Therefore, some component of the DME should be available on a central server. In the following discussion, we refer to this as the **DME server component**. It is likely that the DME server component will be hosted by a Trusted Third Party (TTP) - a PDE service provider.

Since assumption 5 states that PDE sub-networks remain operational even if global connectivity is lost, it can be inferred that DME functionality is distributed throughout the PDE. Hence, **DME components** are distributed within the different sub-networks of the PDE.

In the following, we present the procedure that initialises the PDE, creates the DME server component and fills the mapping table with corresponding information during device registration. The mapping table will contain extensive information on PDE devices, for example, their capabilities, location, current status, reachability, and security parameters. In section 3 we described the notion of home and foreign devices. This distinction is important when considering the details of the registration procedure, which will depend on whether we are dealing with a home or foreign device.

The stages of PDE initialisation presented in this section reflect a more static view of the PDE architecture. In order to accommodate the dynamic aspects, section 6 considers removing a device from the PDE.

## 5.2 DME Server Component

Even if the user has no direct access to their own PDE devices, they will be able to gain access to their PDE via the DME server component at the PDE service provider. The DME server component will contain information describing the current PDE devices. This data will contain the device's identity; time and date of last contact; serving administrative domain, which indicates the identity of the serving network; currently assigned IP address; device resources, and security parameters. The DME server component may also contain information on user location, preferences, and current status.

To access the data on DME server component, all PDE devices must possess the appropriate security parameters required to communicate with it. All PDE devices must also possess some information, locally available, that specifies what the device is authorised to do within the user's PDE. These access and authorisation rights apply to the global PDE, and will be distinct from the device's local access and authorisation rights.

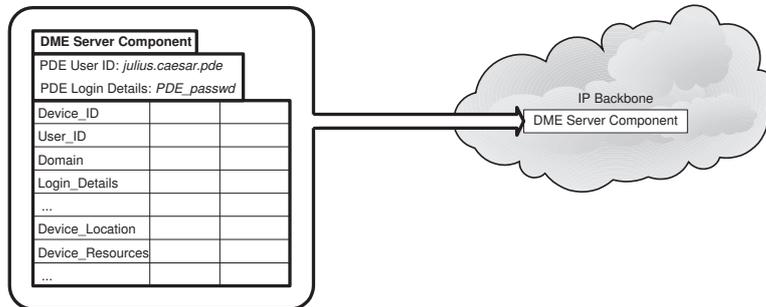


Figure 2: Instantiating the PDE

As we described in section 4, the DME server component and the local DME components cannot be fully synchronised, otherwise the PDE could not be robust. Each *DME component* must therefore contain sufficient information to independently manage its corresponding sub-network. It should also have knowledge of the other sub-networks in the PDE even if some of them are currently unavailable.

The first step for a user to create a PDE is to open a PDE account at a PDE service provider. The user will need to register, and an account will be created in their name. This process may be very similar to creating a mobile phone account and the PDE account may indeed be combined with a user’s mobile phone account. The PDE service provider creates the DME server component for the user. During this procedure the user is provided with a globally unique PDE user ID such as *julius.caesar.pde*, a password, and the network address of the service provider. This allows the user to login to their account at the PDE service provider, and gain access to basic PDE management functionality. The user may at this time specify their preferences and enter their service subscriptions. Alternatively, service subscriptions could be automatically detected when the user adds devices to the PDE. At this point, the user may also receive the corresponding information and software for handling their account.

This completes the first stage in initialising the PDE, but although the DME server component has been created, the PDE does not yet contain any devices. Therefore the mapping table contains no entries. Figure 2 illustrates the creation of a DME server component at a PDE service provider for user *julius.caesar.pde*. In the following, we describe how devices are registered with the PDE to populate this newly created mapping table.

## 5.3 Registering Home Devices with the PDE

In the above, we described how the user creates an empty PDE by opening an account at a PDE service provider, who creates a DME server component with an empty mapping table. We now describe how the user's home devices are added to the PDE, and how information about the device is entered into the mapping table at the DME server component.

### 5.3.1 PDE Device ID

Each PDE device needs to be allocated a globally unique PDE identifier before it can be added to the PDE. We call this identifier the PDE Device ID, or *PDE\_DID*. Since the device must be uniquely identifiable by all PDE service providers, this identifier must be universal. Since the user has already been allocated a globally unique PDE username, *julius.caesar.pde*, each device can simply concatenate its own device ID to the global ID, for example *PDA.julius.caesar.pde* or *laptop.julius.caesar.pde*. This is a naming method that is very extensible and also provides names that can be used in the IETF SIP domain [17] for session set-up.

### 5.3.2 Software Installation and Device Registration

All devices require some basic PDE communications software that allows the device to connect to, and authenticate to, the DME server component. However, not all PDE devices are expected to be IP enabled, and not all devices are expected to have the computational resources required to efficiently implement complex security mechanisms. Devices with such limited resources, or legacy devices, may use a more intelligent device as a proxy. For example, a Bluetooth enabled speaker or display, could use a HiFi or a PC to act as a proxy. The installation of the PDE software on the device or the proxy can be performed automatically or manually.

**Automated installation:** The basic PDE software could be pre-installed, as is the case with Open Mobile Alliance<sup>6</sup> Digital Rights Management software on mobile phones. In such cases, to register the new device, the user only needs to provide the network address of the DME server component, their PDE username, and password. The device can then contact the DME server component via a secure connection, and using the password information, authenticate and register itself. Once registered, the device then uploads information about its features, and

---

<sup>6</sup>[www.openmobilealliance.org](http://www.openmobilealliance.org)

the network contracts that it may possess. In this way the DME can automatically learn of the user's service contracts. For example, the user's UMTS service contract details can be obtained from their mobile phone; their broadcast services contract details can be obtained from their set-top box; and their email account details obtained from their laptop. For devices that do not have direct Internet access, a secondary device that has already authenticated to the DME server component, acts as a proxy authenticator. The proxy device may connect to the new device allowing the user to validate the password and DME address. Thus a stolen device cannot be used to add new devices to the PDE unless the user's password is also compromised.

**Manual installation:** In the case of manual initialisation, the PDE service provider will supply the user with software that they can install on each device. Since these devices are home devices the user has access rights, as administrator or user, and can login to their local device account to install the software. Once the software is installed, the device becomes PDE-aware and has some component for managing the PDE. If the user has no administrator rights, this software may have to be installed by an administrator. The registration procedure can then proceed by contacting and updating the DME server component as described above.

As an alternative to either of the above two procedures, the DME server component may be accessed from any other device via the Internet, and the new device details manually entered.

This completes the first step of the registration process. The new device may now communicate with the DME, and the DME is aware of the new device's capabilities. In the manual approach the user is required to be competent at software installation, which may be an impediment to user acceptance. With the automated approach, the only action user has to perform is to enter their PDE username, password, and service provider network address. From the user perspective, this is not significantly different from the online banking systems currently in use, and should be quite acceptable. The automatic approach also eliminates the problem of users being unable to install software due to problematic hardware/software configurations.

### 5.3.3 Authorisation and Delegation

An important part of the registration process is to define the authorisation requirements for the device. A user's access rights will depend on whether they are logged on directly to the device, or remotely, through the PDE. The

user may also permit others to access the device. The device registration process must, therefore, specify the authorisation requirements and delegate access rights to the device's resources. This can be done by explicitly defining each user, or class of users, who can access the device, and what they can use the device for in their PDE. These rights must be a subset of the primary user's own access rights. One example would be for a user to grant access rights to a set-top box to a friend or neighbour to allow recording television programmes. In this case, the primary user would act as a Trusted Third Party, and the device would be seen by the other user as a potential foreign device to be added to their own PDE.

Rather than set up authorisation and delegation requirements interactively for each device, it is envisaged that default preferences are defined during stage one and stored in a user profile on the DME server component. The default could be for the user to delegate all their access rights on the local device to the DME component, and therefore make this account completely accessible via their PDE. Another possibility is that the user specifies certain applications and data to be accessible to the PDE whereas others, for example, confidential work data in the office environment, are kept out.

#### 5.3.4 Mapping

From a security perspective, the most important step in the device registration process is to define the mapping of device security credentials in the local domain, to their equivalents in the global PDE security domain. By the end of the preceding step, a device will have a unique *PDE\_DID*; be able to run PDE software; and have established access requirements. The device may be able to connect to the DME server and log in, however, it is not part of the PDE until this mapping is complete.

Each device may exist in a different security domain, for example a home or office network. Within each security domain the user will have a different login details. These could be a local user ID and corresponding local security credentials unique to that domain. The mapping process requires the new device to be associated with a specific local security domain and local user, then the user's *local* login details must be mapped to their PDE *global* login details.

There are two scenarios to consider. After logging in to the candidate PDE device and starting the PDE software, the user may either:

1. Map the device *directly* to their PDE using the DME server component at the PDE service provider.

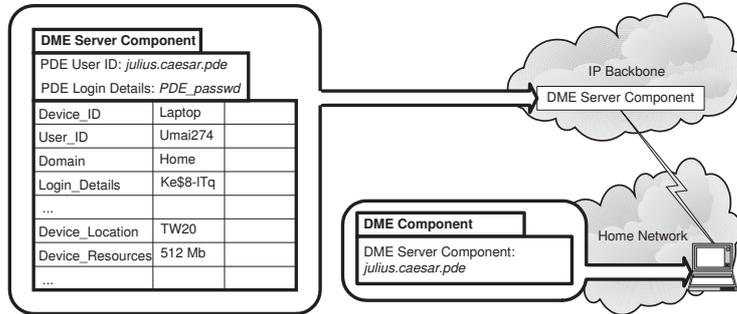


Figure 3: Adding a home device to the PDE

2. Map the device *indirectly* to their PDE using the DME component at another of their PDE devices - the proxy mentioned in section 5.3.2.

**Direct Mapping:** The first case, illustrated by figure 3, represents the default case, where the DME server component at the PDE service provider is directly accessible. The local login credentials of the user need to be mapped onto the PDE login credentials defined in stage one, and the corresponding information entered into the DME server component. To accomplish this, the user has to identify their PDE to the PDE service provider (e.g. *julius.caesar.pde*). A secure tunnel between the new device and the PDE service provider has to be set up, then the user provides a password or other credential to login to their PDE and trigger the registration procedure. This information is sent to the PDE service provider where it can be verified. Next, both sides exchange data relating to the new device, for example, its *PDE\_DID*, how it can be reached, and its capabilities. A protocol is then initiated that derives security parameters for this connection. Therefore, there are two important security aspects to consider, how the secure tunnel is set up between the device and the DME server component to transmit credentials in the first place, and what security protocols are used to derive the security parameters for the new connection. The latter is discussed in [20].

**Mapping by Proxy:** The second case, shown in figure 4, arises when a user wants to add a candidate device to the PDE, but the DME server component is not accessible. In this case the user may register the candidate device via another device in the same sub-network which acts as a proxy. The DME information identified above is then propagated

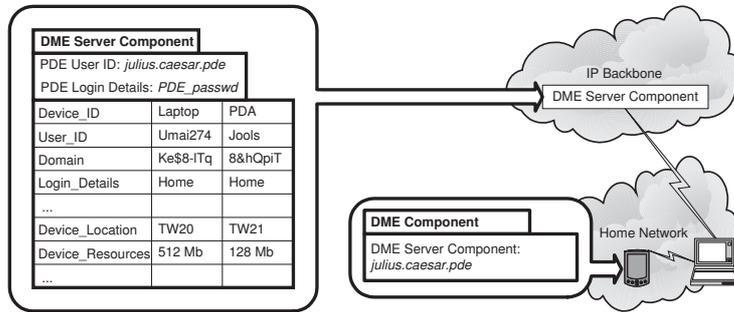


Figure 4: Adding a second home device to the PDE via a proxy

to the DME server component at a later point in time. As was the case with direct mapping, the user logs in to the candidate device to be added and starts the PDE software. The proxy device must be up and running, so that the user can login to the DME component from the candidate device. Again, credentials to authenticate to the PDE must be transmitted through a secure tunnel between the candidate device and the proxy device. Subsequently, a security protocol must be used to derive the security parameters for the new connection. In this case, once registered, the candidate device will become aware the other PDE devices in the sub-network. However it will not be aware of the PDE devices in other sub-networks and on higher layers of the PDE hierarchy until the registration data has been propagated to the DME server component.

In both cases, one more possibility exists to complete the mapping. This is the case where the PDE software is running on the candidate device and that the user logs in *remotely* to the candidate device to perform the mapping procedure.

### 5.3.5 Sub-networks

So far, we have described basic registration procedures for home devices and made reference to the fact that the PDE may be fragmented into sub-networks. We have not yet considered the structuring of the PDE into sub-networks. The concept of a sub-network reflects the lower-level topology of the PDE and specifically the communication protocols at the link layer. The PDE however, is an application layer concept, built upon a heterogeneous and dynamic set of protocols on the lower layers.

The objective of grouping PDE devices into sub-networks arises from the homogeneity of the lower layer security mechanisms in the corresponding environments. Another reason is the easy accessibility of PDE devices within the same environment, for example using WLAN in the home environment, or Bluetooth in the PAN. In terms of the initialisation procedure, it makes sense to group PDE devices into sub-networks according to their lower-layer capabilities. This will have advantages for the management of devices and services during PDE operations. Thus a *PDE sub-network* may be considered to be a fully operational, enclosed entity that can manage itself using a DME component on one of its devices.

## 5.4 Registering Foreign Devices with the PDE

Section 5.3 described the registration procedure for home devices. Foreign devices, by their nature, require a different registration procedure. Foreign devices are not controlled by the user, but by third parties. Ideally these devices should be controlled by a Trusted Third Party (TTP), otherwise the foreign device might compromise the whole of the PDE. A secure procedure is therefore required to allow foreign devices to be added to the user's PDE. This section gives an overview of this procedure.

A precondition for adding a foreign device to a user's PDE is that the owner of that device has already installed PDE software on the device. Thus we may assume that the foreign device is PDE aware and has control over the resources others may access. When requested, the TTP decides whether the user is allowed to add this device to their PDE, and if so, for how long. The right to access the foreign device and add it to the PDE will depend on the relationship between the TTP and the user. We envisage four possible cases.

1. The user has access rights to the device in a specific role (e.g. student of a college, member of a club).
2. The third party knows the user.
3. The user pays for the access.
4. Access is free.

Case 1 means that the user can choose a different identity to their PDE identity to add the foreign device to their PDE. As with home devices, we can adapt Grid computing solutions and apply these to the PDE [20]. Case 2 implies that the user knows the owner of the device, who may be a family

member or friend, who can add the device to their PDE. In both cases the user knows the specific device that they want to add to their PDE. The owner of the device needs to delegate access rights to this user, or to an appropriate class of users. Hence, the third party sets up an account for the user on their device. This effectively turns the foreign device into a home device and the user can use the home registration procedure to add the device to their PDE.

In case 3, the user may be leasing the device from a TTP. Then they get some time-limited access rights which allow them to make the device part of their PDE. In this case, the TTP will have installed the PDE software with corresponding access rights for paying users. In case 4, the device is open to all and does not require payment, similar to anonymous FTP.

Unlike the first two cases, in cases 3 and 4 the user does not know which specific device they want to add to their PDE: the user only requires that the device delivers the desired data or services. In the case 3, the user is willing to pay for this, in the case 4 this is free of charge. Since the user does not know which specific device offers what they need, they have to contact a third party to specify their requirements, make a choice, and add the device to their PDE. This third party services broker may be hosted by the PDE service provider. The broker may present a choice of PDE devices, and the user chooses the device that matches their interests best, for example, which is the most cost-effective. This touches on the Mobile VCE concept of a “digital marketplace” which is described in more detail in [12] and [11]. Having selected a device, the user has to establish a connection between their DME and the third party, over which the registration process takes place. The third party will need to give some security and quality of service guarantees to the user.

Once the connection has been established, there are two different sub-cases to consider, depending on the location of the user and the device. The distinction is between local devices that the user can see, or has physical access to, and remote devices that can only exchange data with the user. The location of the user is therefore an important parameter in resource specification. In the first sub-case the user is remote from the device and must specify the resource, for example, in terms of the delivered service, location, or owner. In the second sub-case, the user can access the foreign device directly and use a simple, manual pairing mechanism such as the Manual Authentication Protocol described in [10].

## 6 Removing a Device from the PDE

Section 5 described the initialisation of the PDE as a two stage process of setting up the PDE framework, and subsequently adding individual devices to this framework. This section briefly describes how devices may be removed from the PDE.

There are several reasons why a device may need to be removed from a PDE. The device could be temporarily unavailable; permanently unavailable (stolen, broken, obsolete); or it could be compromised (security parameters leaked). No action is needed in the first case, when the device will take part in PDE operations at a later point in time. However, if the device has been stolen or is permanently unavailable, then it needs to be removed from the user's PDE. The procedure to remove a device will depend on whether or not the device is accessible.

If the device is accessible to the user, they can log in to their PDE account on the device and trigger a procedure for removing this device from their PDE. This would involve contacting the DME server component and removing the device's *PDE\_DID* from the mapping table. All local confidential PDE data then needs to be removed from the device.

If the device is not accessible to the user, they can still log in to their PDE account at the PDE service provider and remove the device from there. The DME server component removes the device from its mapping table and propagates the information to other DME components as fast as possible.

## 7 Summary, Conclusions, and Future Work

Our objectives in this paper were to describe how a Personal Distributed Environment may be built upon a collection of dynamic heterogeneous sub-networks. We described the motivation for this area of research, and identified other research initiatives that consider purely virtual networks. The major difference between the PDE and more mature technologies is that the PDE is designed to support the services and applications of a single user. The PDE concept takes a user-centric view of consumer devices and merges them into one environment. The dynamic and heterogeneous nature of this environment presents a number of research challenges, and this paper has focussed on the secure initialisation of the PDE.

We described a two stage initialisation procedure that involved the creation of a framework to support the PDE followed by the registration of the user's devices. The creation of the framework introduced the concept of the DME server and local DME components. The registration procedure

required a distinction to be made between the user's home devices and guest foreign devices. In the latter case the registration procedure depends on the relationship between the user and the device. We also described how a device with limited resources may register via a proxy. This is particularly important in the PDE where consumer devices may include low-power short-range devices such as sensors that are not IP enabled and are unlikely to have the computational resources necessary to run complex cryptographic algorithms.

The work presented here represents the foundation upon which we may now construct the detailed security architecture for the PDE. The mechanisms for mapping home devices has been defined in previous work [20], and we are now completing the the procedures to map foreign devices on to the PDE [2].

## Acknowledgements

The work reported in this paper has formed part of the PDE area of the Core 3 Research programme of the Virtual Centre of Excellence in Mobile and Personal Communications, Mobile VCE, [www.mobilevce.com](http://www.mobilevce.com), whose funding support, including that of EPSRC, is gratefully acknowledged. Fully detailed technical reports on this research are available to Industrial Members of Mobile VCE.

## References

- [1] R. C. Atkinson, J. Dunlop, J. Irvine, and S. Vadgama, "The Personal Distributed Environment," in *7th International Symposium on Wireless Personal Multimedia Communications WPMC-04*, Abano Terme, Italy, Sept. 2004.
- [2] K. L. Billington and A. Tomlinson, "Mutual Authentication of B3G devices within Personal Distributed Environments," in *Fifth International Conference on 3G Mobile Communications Technologies, 3G 2004*. London, UK: IEE, Oct. 2004, pp. 452–456.
- [3] Bluetooth, "Specification of the Bluetooth system," Bluetooth, Tech. Rep. v1.2, Nov. 2003. [Online]. Available: [http://www.bluetooth.org/foundry/adopters/document/Bluetooth\\_Core\\_Specification\\_v1.2](http://www.bluetooth.org/foundry/adopters/document/Bluetooth_Core_Specification_v1.2)
- [4] N. Cam-Winget, R. Housley, D. Wagner, and J. Walker, "Security flaws in 802.11 data link protocols," *Communications of the ACM*, vol. 46, no. 5, pp. 35–39, 2003.

- [5] N. Doraswamy and D. Harkins, *IPSec : The new security standard for the Internet, intranets, and virtual private networks*, 2nd ed., ser. Web Infrastructure series. Upper Saddle River, N.J: Prentice Hall, 2003.
- [6] J. Dunlop, R. C. Atkinson, J. Irvine, and D. Pearce, “A Personal Distributed Environment for Future mobile systems,” in *Proc. 12th IST Mobile and Wireless Communications Summit*. Aviero, Portugal: IST, June 2003, pp. 705–709.
- [7] I. Foster, C. Kesselman, J. M. Nick, and S. Tuecke, “The physiology of the Grid - an open grid services architecture for distributed systems integration,” Global Grid Forum, Tech. Rep., June 2002. [Online]. Available: <http://www.globus.org/research/papers/ogsa.pdf>
- [8] I. Foster, C. Kesselman, and S. Tuecke, “The anatomy of the Grid: enabling scalable virtual organization,” *The International Journal of High Performance Computing Applications*, vol. 15, no. 3, pp. 200–222, Aug. 2001. [Online]. Available: <http://www.globus.org/research/papers/anatomy.pdf>
- [9] D. Garlan, D. Siewiorek, A. Smailagic, and P. Steenkiste, “Project Aura: Toward distraction-free pervasive computing,” *IEEE Pervasive Computing*, pp. 22–31, Apr. 2002.
- [10] C. Gehrmann, C. J. Mitchell, and K. Nyberg, “Manual authentication for wireless devices,” *Cryptobytes*, vol. 7, no. 1, pp. 29–37, 2004.
- [11] S. K. Goo, J. M. Irvine, J. Dunlop, A. Tomlinson, and S. Schwiderski-Grosche, “Security Requirements for Mobile Service Provision via a Digital Marketplace (Invited Paper),” in *11th European Wireless Conference, EW '05*, vol. 2, VDE. Nicosia, Cyprus: VDE Verlag, Apr. 2005, pp. 573–581.
- [12] J. Irvine, “Adam Smith Goes Mobile: Managing Services Beyond 3G with the Digital Marketplace (Invited Paper),” in *European Wireless 2002*. Florence, Italy: EUREL, Feb. 2002. [Online]. Available: <http://docenti.ing.unipi.it/ew2002/proceedings/QoS001.pdf>
- [13] ISO/IEC, “Information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements - part 11: Wireless LAN medium access control (MAC) and physical layer (PHY),” International Organization for Standardization (ISO), Geneva, Switzerland, ISO Standard ISO/IEC 8802-11:1999, 1999.

- [14] W. Mohr, “WWRF - The Wireless World Research Forum,” *Electronics and Communication Engineering Journal*, vol. 14, no. 6, pp. 283–291, Dec. 2002.
- [15] —, “The Wireless World Research Forum - WWRF,” *Computer Communications*, vol. 26, pp. 2–10, Jan. 2003.
- [16] P. Pangalos, K. A. Chew, N. Sattari, A. Tomlinson, R. Atkinson, H. Aghvami, and R. Tafazolli, “The Mobile VCE Architecture for the Interworking of Mobile and Broadcast Networks,” in *11th European Wireless Conference, EW '05*, vol. 2, VDE. Nicosia, Cyprus: VDE Verlag, Apr. 2005, pp. 823–828.
- [17] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, “SIP: Session Initiation Protocol,” IETF, RFC 3261, June 2002.
- [18] M. Satyanarayanan, “Pervasive Computing: Vision and Challenges,” *IEEE Personal Communications Magazine*, vol. 8, no. 4, pp. 10–17, Apr. 2001.
- [19] S. Schwiderski-Grosche, A. Tomlinson, S. K. Goo, and J. M. Irvine, “Security Challenges in the Personal Distributed Environment,” in *60th Vehicular Technology Conference, VTC Fall '04*. Los Angeles, USA: IEEE, Sept. 2004.
- [20] A. Tomlinson and S. Schwiderski-Grosche, “Application of Grid Security to Personal Distributed Environments,” in *First International Workshop on Grid Computing and its Application to Data Analysis GADA '04*, ser. LNCS OTM Workshops, R. Meersman, Ed., vol. 3292. Cyprus: Springer-Verlag, Oct. 2004, pp. 68–78.
- [21] WWRF, “The book of visions 2001,” WWRF, IST - WSI Project Version 1.1, Dec. 2001.