

# OpenVMS Technical Journal

## Advanced Server for OpenVMS

Hans Hosang

Advanced Server Competence Center Europe, Netherlands

### Overview

This paper is meant to provide you an insight in Advanced Server for OpenVMS.

Discussed are the usage of the product, like its use in clusters and the roles of the server in various domains depending on the Microsoft functionality that you use.

I will try to shed some light on the fields of troubleshooting, monitoring, performance and tuning.

I will briefly talk about future plans for Advanced Server.

### Introduction

Advanced Server for OpenVMS, a file and print server based on Microsoft Windows NT code, gives you most of the Windows NT functionality on OpenVMS.

Printing is one of the Windows functionalities that are supported by Advanced Server. If you have printer queues on your server, you can setup these printers to be used by your Windows based PCs. In the last version the ease of use of printers has been transformed to the way Windows does it. You can manage your printers from any Windows PC and install printer drivers on the server to be used when you setup this printer on a PC.

Windows functionality like the support of domain structures is natural to Advanced Server. Advanced Server can be leading in a domain, maintaining the user community for an organization. Advanced Server can also play a Backup role in a Windows domain or just be a member server. Supporting domains can also mean supporting inter-domain relations like trusts and we do.

You can separate your user's login domain from resource domains and base your security on the user's login domain by trusting that domain.

As we speak about trusting, OpenVMS still is a widely trusted and reliable operating system.

Many customers use Advanced Server to share out the data of their OpenVMS applications to the world of Windows, either for presentation purposes or to use the data in other programs.

Amongst the largest users of Advanced Server are banks, stock exchanges, lotteries and manufactories.

So basically anywhere where the reliability of OpenVMS is needed you can find Advanced Server.

There are other ways of using Advanced Server; several customers use it as the primary source of security in their Microsoft domain. Here again the reliability of OpenVMS is the main reason for the use of Advanced Server.

Before going into more detail about the functionality, let's view some history.

### History

Formerly known as PCSA and Pathworks, Advanced Server has been around since 1988 and has lived through many changes.

A brief overview:

PCSA ( V3.x and lower ) 1988

Disk services over LAST – LAD (most important)

Most information was stored, not shared, in a container file on OpenVMS. Sharing was only possible if the container was “mounted” read-only.

File service over DECnet

From version 2.0 onwards it was possible to directly share files from the OpenVMS file system. DECnet was the only protocol used in those days.

PATHWORKS V4.x ( MSNET server ) 1991

Remote boot using LAD (Disk Service)

As all you needed to boot a PC fitted on a floppy, you could create floppy container files to boot PCs off the server and thus secure and centrally control the boot environment.

TCP/IP protocol added for File Service.

PATHWORKS V5.x ( LAN Manager 2.2 server ) 1994

RIPL Remote Boot (from File Service)

Slowly abandoning the disk service container files, remote boot was transformed to start off a special file service. This made it more easy to make changes since you didn't need to dismount the read-only container file for every change.

NETBEUI protocol added for File Service

Nowadays only supported on OpenVMS V5.5-2 on VAX ( V5.0F ECO 2 )

PATHWORKS / Advanced Server V6.x 1997

Now based on Windows NT 3.51 code (LAN Manager 3.0)

Based on code from AT&T most Windows NT 3.51 functionality became available on OpenVMS.

No remote boot, no Disk services

Old technologies were discarded, Windows became too big for remote boot, giving network load issues and container files were no longer used as file services became faster and more common.

Advanced Server V7.2 1998

Introduction of the OpenVMS registry

Functional this release was the same as Pathworks V6.0. The only differences were the replacement of an ini file by the registry and the introduction of ODS5 disk structures with the use of extended character set and lowercase filenames.

Advanced Server V7.3 2001

Member server function

As it lacked from the AT&T code Member server functionality was added by our own labs to enable us to participate in Windows 2000 Active Directory domains. This functionality is also available for VAX in Pathworks version 6.1

Advanced Server V7.3A (ECO-2) 2002

Windows NT 4.0 based server

Lots of bug fixes and performance improvements over version 7.3.

This document is based on the functionality of the latest version: V7.3A patch level ECO2.

This version is a complete rewrite of large parts of the product to make the transition from Windows 3.51 to Windows 4.0 code base. The functional changes that have been made in Service pack 4 of NT 4.0 are all implemented in this release. This makes the product interact and function well within Windows 2000 and Windows 2003 Active Directory domains.

## Domain Functionality

In a Windows NT domain there are three different roles that a server can perform within the domain.

### 1. **Primary Domain Controller (PDC)**

This is the domain master in maintaining the security identifiers in the domain. All users and groups are given a security identifier by the Primary DC to be used for assigning access throughout the domain. The Primary DC distributes all changes in the domain to the Backup DC's. Without the Primary DC, no change in the domain, like a user's password, is possible.

### 2. **Backup Domain Controller (BDC)**

A Backup DC has two main functions. First it aids the Primary DC with logon validation. All users that logon to the domain are checked, validated, by a Domain Controller, PDC or BDC. This is a form of load balancing. A second function of a Backup DC is to maintain domain consistency and availability. Should your Primary DC disappear in smoke, you still have your domain based on the domain copy that is held by all Backup DC's. In such a disastrous event you can promote a Backup DC to Primary DC in order to allow changes to the domain.

Note: Running a domain without a Backup DC is like living dangerously, it's not the question IF you will hurt yourself but when. Since all Windows machines change their password autonomously, you can never be sure that a restore from backup will bring back your domain completely. Some PC's or servers may have to be brought back into the domain manually because they changed their password after the last backup. Also trusts to other domains may fail the same way.

### 3. **Member Server**

A Member Server is only file and/or print server, it does not maintain a copy of the domain database and requires the availability of a Domain Controller for validation to allow users access to its shares.

Advanced Server for OpenVMS V7.3A can perform each of these three roles.

Each of these roles has its specific usage and the list below is what we see most in the field (numbering as above).

1. Advanced Server as Primary DC is most often used in environments where system management is split up and where OpenVMS and Advanced Server are managed separately from the Windows environment. In these cases a trust is needed between the Advanced Server world and the Windows world to make management easy and security common. See the next section for more information about trusts.
2. As in any domain, a Windows NT or Advanced Server domain needs a Backup DC to secure the validity of the domain user database (SAM). This BDC can be any out-of-the-box PC running Windows NT or another OpenVMS or UNIX machine running Advanced Server.

Advanced Server can also be a BDC in a Windows 2000 domain as long as this domain is running in mixed mode.

3. The member role for Advanced Server is most often used in conjunction with a Windows 2000 or Windows 2003 domain. If such a domain is running in Active Directory integrated mode, this is the only role which Advanced Server can perform in such a domain.

As we briefly touched inter-domain relations above, I will now go into more details on relations between domains.

### **Inter-Domain Relations, Trusts**

With many customers, there are multiple domains on the network. This can be due to various departments managing their own domains, a geographical separation or just a boundary between production and test environments.

Often it is needed to give users from one domain access to resources in another domain. In the past this had to be done by creating a username in both domains for such a user and give him the task to keep the passwords in sync. With Windows NT, and Pathworks version 6, the option was introduced to have a username in only one domain and use it throughout the network. To do this, the other domains have to trust your login domain. Once a trust is in place, you can take a username or a global group from the trusted domain and use it in the security masks of trusting domains. Please note that a domain local group cannot be used beyond the own domain. A domain local group can also not be used on a member server because a member server has its own local groups.

Like any Windows domain, a domain with Advanced Server in it can trust any other domain or can be trusted. There is a configurable maximum to the number of domains that can be trusted. If you need to configure your domain for more than 31 trusts, increase the following registry setting;

```
Key: SYSTEM\CurrentControlSet\Services\AdvancedServer\ProcessParameters  
NumClient_Session REG_DWORD 5 - 128.
```

Default value is 32.

Description:

Limits the number of trust relationships that a server can maintain with other domains.

This value should be at least one greater than the number of domains trusted by the servers' domain.

You can best set such a registry value by the REGUTL command on OpenVMS to prevent typos, since REGUTL knows which registry keys are known to the server.

For example:

```
$ REGUTL SET VALUE * NUMCLIENT_SESSION 44
```

There is a server configuration element to take into account when using trusts.

Trusts use up sessions, so you will have to increase the number of PC clients that you configure your server for if you have a large number of trusts.

For details about configuring your server, see the section entitled: Monitoring and Performance.

### **Protocols**

For PCs and servers to communicate with each other you need a common language or protocol.

There are several protocols available for machines to communicate, including TCP/IP, AppleTalk, IPX, DECnet and many more.

Advanced Server can be configured to use one or more three protocols:

- DECnet
- NETBEUI
- TCP/IP

Because the TCP/IP protocol is used more and more these days, we normally see systems that are configured to use TCP/IP only. Some details about this protocol; TCP/IP is built around the usage of network numbers or subnets. This feature is mostly used for the purpose of routing the protocol between several sites of a company but can also be used to divide the local network into manageable parts. Apart from the necessary segmentation of your network, the principle of subnets can give you some headaches. Windows NT connectivity is partly based on the principle of browsing and this functionality is not routable between subnets or sites. There are many OpenVMS systems using multiple TCP/IP subnets on multiple network cards. This could bring a challenge for Advanced Server, especially when this happens in a cluster. For Advanced Server it is not supported to have individual cluster nodes in different subnets only. They must have, at least, one subnet in common on all nodes of the cluster. This has to do with the way Microsoft designed browsing.

Since the NETBEUI protocol and some functionality of the other protocols, like browsing are not routable you may have to configure Advanced Server to use a specific interface on a system that has multiple interfaces.

This can be done by defining a logical name per protocol that you use.

```
For DECnet:    $ DEFINE /SYSTEM/EXEC NETBIOS$DEVICE          FWA0 :
For NETBEUI:  $ DEFINE /SYSTEM/EXEC PWRK$NBDAEMON_DEVICE    FWA0 :
For TCP/IP:   $ DEFINE /SYSTEM/EXEC PWRK$KNBDAEMON_DEVICE    FWA0 :
```

When you define the above logical name for TCP/IP you also have to specify the IP address for this device with the following logical name:

```
$ DEFINE /SYSTEM/EXEC PWRK$KNBDAEMON_IPADDR "16.198.227.15"
```

You can also use this method if you have an interface that is not yet recognized by Advanced Server.

Should you be unable to have at least one subnet in common on your cluster nodes, you will encounter problems with the browser service like a permanent state of "Start Pending" on some nodes in the cluster. In such a case you will have to disable the Browser service. Do this by setting a registry key:

```
$ REGUTL SET PARAM * MAINTAINSERVERLIST NO /CREATE
```

After a restart of Advanced Server you will no longer see the PWRK\$LMBROWSER process and you will get a report of this in the system event log, where it reports you the message numbers 7024 and 2550, indicating that the browser service had nothing to do and exited at startup.

So far for protocols, now I come to the way a machine can find another one on the net.

## **Name Resolution**

There are various ways to file a combination of name and address for a machine. You don't want to define every machine on every PC, this can largely be automated.

Depending on the version of operating system there is a variety of options that are used in different sequences to find an address for a machine on the network.

As in Windows NT, name resolution for Advanced Server is based on a combination of broadcasting, WINS and DNS. WINS, (Windows Internet Naming System) is for Windows NT, and thus for Advanced Server, the preferred way to find an address for a machine. This is also the first option to try on the network, after checking the local memory cache. If WINS doesn't provide an address for the machine you are looking for, Advanced Server will use broadcasting to get the address. Broadcasting, however, will only travel as far as the TCP/IP subnet reaches. This makes broadcasting very limited and not so useful. Broadcasting relies on the machine itself to answer the call so it will only be effective if the machine is in your own subnet. As a last resort, DNS can be used, in order to find an address for a machine name.

I did not mention a domain name in the sequence above. This is because for searching a domain name, DNS will NOT be used. In the search for a domain name or a domain controller, Advanced Server uses WINS and broadcasting only, Just like Windows NT. This makes it important to keep WINS if you are transferring your domain to a Windows Active Directory domain, based on DNS.

The above mentioned local name cache can be preloaded by the use of a file called: PWRK\$LANMAN:LMHOSTS. (without extension). This file has the same layout as in Windows;

```
<TCP-IP Address> <hostname> [#PRE [#DOM:Domain-name]]
```

For example:

```
16.11.231.18      TESTPDC          #PRE #DOM:TESTDOM
```

This line loads the name TESTPDC in cache with the address 16.11.231.18.

This line also loads the domain name TESTDOM in cache and associates it with the same address.

This way your machine knows that it has to go to 16.11.231.18 if it needs to contact a domain controller for the domain TESTDOM.

You can configure the use of WINS, DNS and the use of an LMHOSTS file using the configuration utility: \$ADMIN /CONFIG, in the TRANSPORTS section. By default none of them is enabled and your system relies on broadcasting only, something you don't want to do in a modern domain.

New in version 7.3A-ECO2 is that you can edit and reload the LMHOSTS file without restarting the server. To do this, you can use the command \$ NBSHOW KNBCACHE RELOAD. The command \$ NBSHOW KNBCACHE, without the reload option, can be used to display the current content of the name cache.

## Troubleshooting name resolution

Troubleshooting name resolution can best be started using the following command:

```
$ NBSHOW KNBSTATUS <NAME>
```

You should expect to receive a list of claimed names from the server or client requested, like the output from the Windows following command:

```
$ NBTSTAT -a <NAME>
```

Following is sample output (partial) of \$ NBSHOW KNBSTATUS command with explanations of the various claimed names:

```
Pwop01-System > nbshow knbstatus pwop08
```

```
.....
```

```
Local name table (11 names):
```

Name	Soc Num	Status	explanation
------	---------	--------	-------------

PWOP08	x20	1	Unique	Registered	<- Server name, server service
PWOP08	x00	2	Unique	Registered	<- Server name, workstation service
PWDOM2	x00	3	Group	Registered	<- Domain name
PWDOM2	x1c	4	Group	Registered	<- Domain Controller in PWDOM2
PWDOM2	x1e	5	Group	Registered	<- Used for Browser elections
PWOP08#D	x00	6	Unique	Registered	<- Internal usage
PWOP08#B	x00	7	Unique	Registered	<- Internal usage
PWDOM2	x1d	8	Unique	Registered	<- Master Browser
^^_MSBROWSE^^	x01	9	Group	Registered	<- Master Browser
PWDOM2	x1b	10	Unique	Registered	<- Domain Master Browser (PDC)
PWOP08_65	x00	11	Unique	Registered	<- Internal usage

Important in this output are the server name and the domain name. They are both listed multiple times in this output and here you can see the functions this server is performing. You can also use this method to search for the PDC or PDC emulator (Windows 2000) of a certain domain. If you encounter difficulties in joining a domain during the configuration of Advanced Server you can use the command `$ NBSHOW KNBSTATUS` to see if you can resolve the domain name to find the PDC. Since version 7.3A, the command `$ NBSHOW KNBSTATUS <NAME>` accepts a third parameter where you can specify the last byte of the NetBIOS name (listed under the column “soc” in the table above). For example, if I issue the command `$ NBSHOW KNBSTATUS PWDOM2 1B`, I will get the same output as above while searching for the PDC in domain PWDOM2.

You can use this if you encounter problems while joining a domain during initial configuration. You can do this because the `PWRK$KNBDAEMON` process that interfaces to TCP/IP is not shutdown when you leave the configuration procedure.

So far for communication between machines, let’s take a closer look at configurations.

## Clusters

### The cluster is the server

This means that there must be one single entity that is seen as the cluster from the outside world.

When you use an OpenVMS cluster, you will have to give this cluster a cluster alias name. You can configure many nodes in a cluster to run Advanced Server. They will all listen to the same cluster alias name.

There is, however, no need to configure all of the cluster nodes for the use of Advanced Server.

### File access in a cluster

It is important to mount the disks that you want to share out to the Windows world on all nodes that you configure to run Advanced Server. Failing to do so will make a file unavailable through some nodes and this will look inconsistent from a user’s perspective. Unlike a Windows cluster, you will be able to access all disks through every node in the cluster and at the same time. This feature, unique to OpenVMS and Tru-64 UNIX clusters, does provide some overhead in the server which may cost some performance.

Should one of the nodes in an Advanced Server cluster leave the cluster, the clients connected to it can resume work through one of the other, remaining, nodes without user intervention. To use this feature, clients must use the cluster alias to connect to the cluster and not a specific node name.

Advanced Server has a set of databases to store the reference to shares, users and parts of the security.

In an Advanced Server cluster there is only one set of these database files.

These files must reside on a common disk that is mounted on all nodes that run Advanced Server. This disk must preferably not be local inside one of the nodes but in a separate storage enclosure to prevent unavailability if one node fails.

### The cluster in a domain

As I told you above, the cluster can and should be addressed by its cluster alias name. This is one name for one or more machines. In the Domain, only the cluster alias name must be registered as a server or domain controller, depending on its role in the domain. It is not allowed to create a computer account in the domain for individual nodes of the cluster. This and the fact that there is one set of databases, mentioned above, means that the whole cluster performs the same domain role. If you make your cluster the PDC, this will lead to the situation that you will see multiple PDC's when browsing the domain. This is intended behavior since the browser gathers information on both the cluster alias as well as the individual nodes. If you look at domain members only, you will just see the cluster alias, not the individual node names; they only exist in Browser information.

In the OpenVMS administrators utility this looks as follows:

```
$ ADMIN SHOW COMPUTERS
```

```
Computers in domain "PWOP":
```

Computer	Type	Description
[PD] PWOP01	OpenVMS (NT 4.0) Primary	Advanced Server V7.3A for OpenVMS
[PD] PWOP02	OpenVMS (NT 4.0) Primary	Advanced Server V7.3A for OpenVMS
[PD] PWOPCLU (Alias)	OpenVMS (NT 4.0) Primary	Advanced Server V7.3A for OpenVMS

Total of 3 computers

```
$ ADMIN SHOW COMPUTERS/TYPE=DOMAIN
```

```
Computers in domain "PWOP":
```

Computer	Type	Description
[PD] PWOPCLU	Windows NT Primary	

Total of 1 computer

Note that when you look at the domain members only, the server comment (description) is omitted since this is browser information and the browser information is not used with this command.

### Dos and Don'ts

This section provides recommendations for Advanced Server and things that you must not do with Advanced Server.

DO:

- A PC client will have more than one session to a server, most likely two sessions per PC. This means that you have to configure our server for a little more than twice the number of PC clients that you expect to use the server.
- You can configure any node in the cluster for as many clients as you expect to use it but make sure to have enough room to accept clients that failover from a node that exits the

cluster. For example, if you have a two-node cluster and 400 clients, configure both nodes for at least 850 sessions.

- All cluster nodes must be in the same TCP/IP subnet.
- Configure the server to use WINS for name resolution.
- Add a multihomed entry for the cluster alias in the WINS server database, even if your cluster has only one node running Advanced Server.

**DON'T:**

- Don't configure any server for less than 40 PC clients. It will also use client sessions for inter domain sessions, trusted domains and browsing.
- All nodes in the cluster **MUST** speak the same protocols. Don't let one speak TCPIP only and the other TCP/IP and DECnet.
- Don't configure your server to speak DECnet if your PDC doesn't. The Primary Domain Controller must speak all the protocols that are in use throughout the domain. See below for details.
- Don't build a domain with only a PDC. Always add at least one BDC to the domain.
- Don't create a computer account in the domain for the individual cluster nodes, just for the cluster alias.

A little explanation on the protocols of your PDC, mentioned above:

Suppose your PDC is a Windows machine that is configured to run TCP/IP only and you want some old PC's to connect to Advanced Server through DECnet.

In this scenario your Advanced Server for OpenVMS could be the only server that speaks DECnet. If that is the case, it will become Domain Master Browser on the DECnet protocol since there is no higher machine that claims this role. The Browser service is not built to have a different role on different protocols so this will bring your PDC in trouble because the PDC should be the Domain Master Browser.

This situation will create an error condition on TCP/IP which is hard to find since the cause is not on TCP/IP but on the DECnet protocol.

You can avoid this by installing DECnet on the PDC. DECnet is supported on all current Windows versions and can be installed from the Pathworks 32 CD kit.

## Licensing

Any connection to Advanced Server must be licensed.

Nowadays there is only one type of license that you need to install in order to get access to shares on Advanced Server. This is the Client Access license, named PWLMXXXCA07.03.

There are two principal ways to use this license.

1. Server-based licensing

If you have only a few servers you can buy a set of licenses for each node and just load the license into LMF. In this setup, you configure Advanced Server to NOT use the License Server.

If you do this, Advanced Server will subtract a license for each PC that makes a connection to the server. There is nothing to manage with this license setup, just make sure that you have enough licenses on each server.

You can check the license usage with the command `$ PWLIC`. Please keep in mind that

license usage in a cluster is always cluster-wide, so PWLIC will show you the combined license usage for the whole cluster.

One more remark on licenses in a cluster: The license processes on the cluster nodes talk to each other to give this combined license usage. You **must** setup your cluster to use **one** license database for the whole cluster. If you don't do this and use multiple license databases (e.g. per node), Advanced Server nodes will negotiate about the total number of licenses for the cluster and the node with the smallest number of licenses will overrule the others. To have multiple license databases in a cluster is officially unsupported.

This is the easiest and preferred setup and is called server-based licensing.

## 2. Client-based licensing

If you have many servers or when PCs make use of multiple servers, it may be cheaper to setup one or more servers as a license server.

To do this, like with option 1, load your licenses into LMF on one node in your network and configure Advanced Server on this node to run WITH License Server. Now the License Server will allocate all these licenses and will give them to PCs that request a license.

To make a PC client request a license, you need to install license software on each PC client that needs to connect shares on Advanced Server. This license software is available on the Pathworks 32 CD kit and on the share pwlicense on each Advanced Server node.

With the license software installed, the PC client will present its license to any server it needs to make a connection to. The server that a PC client with a loaded license connects to does not need to have any license installed.

This license will remain with the PC client that requested it until it is released manually.

The license server will maintain a database for all the licenses it has given out.

You can create reports on the usage of licenses through the license utility:

```
$ ADMIN /LICENSE
```

As mentioned above, you need manual intervention to release a license from the database if a PC client leaves the network.

This is called Client-based licensing and is more labor intensive.

## 3. You can also create a mix of options 1 and 2.

If you run your server with the License Server, you can move a number of licenses to a special group called "Server-based". By this mean you can make a new PC client connect to this server to install the client license software from a share on this server.

You can also install the client license software from the Pathworks 32 CD kit that comes with your OpenVMS distribution.

License shortage will be reported in the general event log that you can examine with the following command:

```
$ ADMIN /ANALYZE [ /SINCE [= <DATE-TIME> ] ]
```

## Monitoring and Performance

Basic monitoring of Advanced Server is pretty easy and can be done at three places.

1. There is a general error log about OpenVMS related issues and licensing.  
This file is: `PWRK$ROOT:[000000]EVTLOG.DAT`.

You can access this error log with the following command:

```
$ ADMIN /ANALYZE [/SINCE[=<TIME>]]
```

There is no need for the server to be running when you issue this command.

All you should see here is a few startup messages and maybe an autoshare message at startup time when you have disks with a very long label, like a mounted CD.

It's a good practice to cleanup the error log at a regular basis by the command:

```
$ ADMIN /ANALYZE /PURGE [/BEFORE=<TIME>]
```

2. As a second source of information, there is the Windows NT compatible event log.  
These files are: `PWRK$LMLOGS:SYSEVENT.EVT`, `APPEVENT.EVT` and `SECEVENT.EVT`.

You can access these event logs with the following command:

```
$ ADMIN SHOW EVENT [/FULL] [/TYPE=SECURITY | SYSTEM | APPLICATION]
```

The system event log is default when `/TYPE` is omitted.

You can also examine these events through the event viewer application on your Windows platform.

It's also good practice to cleanup this event log at a regular basis by the command:

```
$ ADMIN CLEAR EVENTS [[/TYPE=SECURITY | SYSTEM | APPLICATION]
```

Alternatively, you can use the Windows event viewer.

You should make sure that this event log is setup to overwrite older events as needed or you will lose messages. This setting can only be changed from the Windows platform.

3. The individual processes that makeup the server produce log files.

There are two directories where you can find these log files. These are: `PWRK$LMLOGS` for those processes whose name starts with `PWRK$LM`, like `PWRK$LMSRV`, the other location is `PWRK$LOGS` for processes whose name starts with `PWRK$`, without the `LM`, like `PWRK$KNBDAEMON`.

The most important of these log files is the actual file server log file,  
`PWRK$LMLOGS:PWRK$LMSRV_<NODE NAME>.LOG`

These filenames always contain your machine name to distinguish between nodes in a cluster.

A simple `TYPE /TAIL` command will show you the end of the file to take a quick look at the current status.

We consider it good management if you take a regular look at these three places for errors.

More monitoring options will be discussed in the next section.

## Performance and Tuning

As with any product, you have to know what to expect if you want to do some tuning.

This means that you will have to know what the normal workload is for your system and for Advanced Server, a baseline. In other words, if you take your first look at performance when your users start complaining about the performance you don't know what you are looking at.

Advanced Server performance breaks down into OpenVMS performance and Advanced Server itself.

Advanced Server performance should be about equal to a Windows NT server.

Only when many small files are involved, like in a user profile, we are a bit slower.

BUT, as often done, you should not compare last century's Alpha server 2100 with a Proliant server that you bought recently, even if they are about the same size.

My intention with this chapter is to give you a few handles to do basic tuning to your server, not to give a detailed, five day training in a nutshell.

What aspects are there in tuning?

There are CPU, memory, disks, the transport protocol and the network itself.

Then there is OpenVMS and the Advanced Server product.

Let's start with Advanced Server itself.

## Configuration

When you have installed Advanced Server, you have to configure it. In the first part of the configuration you will go through a menu that you can also start at a later stage with the command \$ ADMIN /CONFIGURE.

The first screen of this menu contains important settings for your server.

### Example

```
+----- Advanced Server Configuration for node PWOP01 -----+
| Options  Help |
|               |
| +Server's Client Capacity-----+ |
| |               | |
| | ( ) Maximize Client Capacity Using AUTOGEN/Reboot | |
| | ( ) Maximize Client Capacity Without AUTOGEN or Reboot | |
| | (*) User Supplied Client Capacity | |
| |               | |
| | Client Capacity: 50 | |
| +-----+ |
|               |
| Percent of Physical Memory Used: 80 |
| Data Cache Size (Kbytes): 131072 |
|               |
| Maximum Concurrent Signons: 10 |
| OpenVMS Process Priority: 9 |
|               |
| +-----+ +-----+ +-----+ +-----+ |
| | Verify | | Quit | | Advanced... | | Transports... | |
| +-----+ +-----+ +-----+ +-----+ |
|               |
```

+ Test for supportable configuration —————+

Some remarks about this first page, in order of importance:

- Never use a client capacity below 40, your server normally needs sessions to other servers and domains. And each PC client creates, most often, two sessions to your server. So, at average, you setup a server for 2 times the number of expected PCs plus some extra (at least 20 extra).
- Choose a proper value for your data cache size. Shown here is the upper limit which could be an overkill for a 15 user system like this one (Client Capacity: 50). 32000 to 64000 is often adequate. The default value of 8192Kb can be considered a small cache size but could do for a small system in a small domain. The Advanced Server data cache only caches open files! I'll show you in more detail how to examine the data cache, later in this article.
- Percent of Physical memory used: 80. This is the default and this number is only used to make a calculation for the system parameter WSMAX (process memory usage) and to see if your system can support the number of clients that you specify. It is good practice to reduce this number only if there is also a lot of interactive usage on your system. Please keep in mind that your server process may be representing a 1000 user workload or more and in such a case, the server will need quite an amount of resources.

There are two other sections in this menu:

- The Advanced section of the menu is normally not changed.

Go here if you need to change permission settings to include OpenVMS rights like the usage of resource identifiers in your security masks.

Or, to turn off the Open File Cache. This is a file header cache only, not the file data and could be in your way if you have an OpenVMS job that needs to process a file as soon as it is put on the server. If the OpenVMS job can wait the five seconds of the Open File Cache timeout, then do that and leave the Open File Cache enabled. The Open File Cache is a major win for batch jobs and applications like Word that open and close a file 10 or more times before they really start to read it.

- The Transports section of the menu.

Normally you do visit this part of the menu to turn off DECnet and enable TCP/IP. You must choose a way of name resolution like WINS or DNS. Remember that Advanced Server uses WINS as the primary way for name resolution, so if you need to access any machine outside your TCP/IP subnet, you must have a WINS server and enable WINS.

We do advise you, these days, to enable the usage of the LMHOSTS file, even if it is not there, since in this version you can reload it dynamically. This can be very handy should you be confronted with sudden changes in your network. The way to reload the LMHOSTS file is through the following command:

```
$ NBSHOW KNBCACHE RELOAD
```

## Configuration When the Server Is Running

You can use the `$ ADMIN /CONFIGURE` command at any time during normal system operation; it will not interrupt the system or the server. One good moment to use this command is when you anticipate growth of your PC community. You can use it to see if your sever needs parameter maintenance in order to accept more pc client connections. Increase the client capacity and do "verify" to see what message it will produce. You can always bail out of this action by not accepting the change.

My advice: don't let the menu run AUTOGEN for you because it will never use the feedback option.

## Gathering Data

As I wrote before, when you want to do tuning, you need to know what the system's normal behavior is.

With the command `$ PWMON` you can get a good overview of server activity.

**Notes:** In the older versions of Advanced Server you have to set your screen to a width of 132 columns to use the PWMON command.

All PWMON commands in version 7.3A-eco2 will work on a screen that is set to a width of 80 columns but several commands will give you more information if you set your screen to 132 columns.

I'll give a few examples of how to use PWMON.

Let's start with the data cache, as promised.

I'll show you the most important part of the output to keep things limited.

```
$ PWMON DATA_CACHE
```

```
Advanced Server Data Cache/Process=PWRK$LMSRV

Cache buffer size.          8192 Cache buffer count          8192
Cache buffers free         7808 Cache buffers used            384
Nr of assigned handles     22  Cache user limit              0

Data lookup rate          5423213 Cache misses                0
Cache hits                 4298693 Cache hitrate (%)           79
```

In the above screen, first line, you see a count of 8192 buffers, all 8192 bytes in size. This system is configured for 65536 Kb data cache size with `$ ADMIN /CONFIG`. (See for a reference the screen shot of the menu two pages back where you see "Data Cache Size (Kbytes): 131072").

On the second line you see there are 7808 buffers free which means there is plenty of free space.

On the bottom line you see a cache hit rate of 79%. Depending on your system usage, this can be good or bad.

In this case the usage of the system consist for 90% out of large files that are opened and closed regularly and both read and written to, so the hit rate number is good. Remember only open files are cached, so if your application keeps the file open, cache results will improve.

## Another Scenario

Your system is generally OK but sometimes very slow.

In this scenario it is likely that someone is hammering your server with some batch job, you can find out quite simple who that is by the command: `$ PWMON CLIENT`

If you issue this command without parameters or switches you will just get a list of connected PCs and how many commands they issues to your server. See below:

Advanced Server Clients

Name	Tot. SMB's	Tot. API's	Tot. RPC's	Tot. Alrts
CLIENT_LM_BARK01	31	0	0	0
CLIENT_LM_HOUBENE02	1015	0	0	0
CLIENT_LM_HHOSANG02	1468	0	0	0
CLIENT_LM_VANGEESTR03	113765	0	0	0
CLIENT_LM_UTOCWARN1	14447	0	0	19
.....				

This output has little value if the number of PC's connected to your server is more then will fit on one screen. You will scroll from page to page without getting the real overview.

It will be much handier to make use of the new qualifiers available in this version. There is /OPERATIONS, /RESOURCES and /TIMES but also /TOP=(OPERATIONS | RESOURCES | TIMES) to sort on the usage by the PC clients.

PWMON client is one of the PWMON commands where you will get more columns at a screen width of 132.

```
$ PWMON CLIENTS /TOP=OPERATIONS
```

Advanced Server Clients by top operations						
Name	File Opens	File Locks	Dir Lookup	Transact's	Latest SMB	Latest API
CLIENT_LM_VANVELZEN01	26	12	0	199	5	
CLIENT_LM_UTOJSCHO1	2	0	6	53	113	
CLIENT_LM_HOUBENE02	43	40	15	652	113	
CLIENT_LM_HHOSANG02	46	26	11	556	50	

I omitted the last column in the display to make it fit on the page; this is showing the latest API call.

This screen will give you an overview of the top working PCs. Please pay special attention to the column "Dir Lookup" as this is known to generate a large load on the server and could point you to a badly designed program, running on a client.

Should the badly designed application be unavoidable, your next step is to take a look at the disks; how is the server dealing with disk caches?

**Disk Cache Statistics**

```
$ PWMON ODS2
```

This will show you disk related caches in Advanced Server.

The important ones are File ID cache, Directory cache and the Path cache. All of these caches should have a hit rate of around 90% or more.

Advanced Server ODS2 Cache/Process=PWRK\$LMSRV

Nr of PATH cache lookups	1246127	Nr of PATH cache hits	1222885
Nr of PATH cache misses	23242	<b>PATH cache HIT RATE (%)</b>	98
Nr of PATH cache invalidates	0	Nr of PATH cache searches	0
Nr of DIR cache lookups	2078357	Nr of DIR cache hits	1962628
Nr of DIR cache misses	115729	<b>DIR cache HIT RATE (%)</b>	94
Nr of DIR cache invalidates	1	Nr of DIR cache refreshes	614

Nr of DIR	cache thrashes	3553	Nr of DIR	cache CTEXT's free	256
Nr of DIR	cache entries free	3	Nr of DIR	cache blocks free	1391
Nr of DIR	cache buf's total	4	Nr of DIR	cache buf's in use	0
Nr of FID	cache lookups	5842722	Nr of FID	cache hits	5452852
Nr of FID	cache misses	389870	<b>FID cache HIT RATE (%)</b>		93
Nr of FID	cache invalidates	35963	Nr of FID	cache searches	13334112
Nr of FID	cache thrashes	0	FID	cache size	1152519
Nr of FID	cache entries	2250			
Nr of DAT	cache lookups	28588	Nr of DAT	cache hits	26400
Nr of DAT	cache misses	2188	DAT	cache HIT RATE (%)	92

Most important in this page is the FID cache. As all Microsoft products want to know information from the file header like date and size, this cache is used a lot.

To influence the FID cache you can run the configuration utility `$ ADMIN /CONFIG`, go to the advanced menu and increase the number of open files per client. The number of FID cache entries can grow up to 16384.

A restart of the server is needed to effectuate this change.

If you have a system that is configured for a small number of clients, the above modification will not lead to much of a change. In this case you can better create a logical name as follows:

```
$ DEFINE /SYSTEM /EXEC PWRK$TEMP_ODS2_FID_CACHE_SIZE 16384
```

And then restart the server.

Defining this logical is, of course, not a one time event. You must put this in the system startup procedure.

To check the value that is used by the server you can check the file server log file;

```
$ SEARCH PWRK$LMSLOGS:PWRK$LMSRV_UTRACK.LOG FID_CACHE
21-APR-2004 13:09:24.45 20400458:007B3948          ODS2_FID_CACHE_SIZE:    2250
```

This shows that our server is using the value of 2250 for the file ID cache.

The same mechanism can be used for the other ODS2 caches.

**Note:** Everywhere where ODS2 is mentioned, this is also used for ODS5 disks.

## CPU Usage

I've recently come across a customer who complained about a very high CPU usage.

After a lot of testing I ran the command `$ PWMON CLIENT /TOP=OPERATIONS` as described above.

I discovered that a few PCs had high and rapid increasing numbers in the transactions column.

Transactions are commands that are not file related like domain queries or device queries.

To find out what such a PC is doing you have to take a look at the network.

You can take a snapshot of the traffic between a PC and the server using the command `$ TCPTRACE`.

TCPTRACE is a part of HP TCP/IP for OpenVMS and is not available if you use TCPware.

The syntax is:

```
$ TCPTRACE /PACKET=10000 /OUTPUT=FILE.TXT <IP-ADDRESS-OF-PC>
```

When you omit the TCP/IP address of the PC you will capture all traffic to and from the server.

When looking at the output of `$ TCPTRACE`, I found that the PCs were constantly looking for printer information while they had no need for a printer on this server. I de-installed this printer on those clients and the amount of CPU usage decreased by about 5% per PC.

The easiest way to take a look at the output of `TCPTRACE` is through the use of the product `ETHEREAL` which you can download for free from <http://www.ethereal.com>. But in this case all the traffic to this PC contained the comment of the printer share, a type of the output file was sufficient. This example shows that you don't always need in depth knowledge of network protocols to see what is happening. Just don't be scared to take a look.

Another known CPU intensive item is large directories. If you have directories that contain thousands of files then you will certainly have a high CPU usage if they are accessed frequently. The simple reason is the fact that Windows requests information from the file header of each file. The only thing you can do against this is to reconsider the layout of your directory tree.

## Troubleshooting Tips

When something "undefined" goes wrong, where do you start searching?

Say, a user calls with the report that your server is unresponsive.

I'll give you a few hints and tips where to look.

To start, make sure that you defined the specific commands for Advanced Server in your `LOGIN.COM` by:

```
@SYS$STARTUP:PWRK$DEFINE_COMMANDS.COM
```

I'll start with the command; `$ PWSHOW [CLUSTER]`

This will give you an overview of all Advanced Server related processes, per node or cluster wide, if you specify the parameter "cluster". You have to know what processes your server normally runs to discover if one is missing. Normally you should see between 6 and 10 processes.

The minimal process list is:

20400456	PWRK\$ADMIN_0	LEF	6	56	0 00:00:00.01	153	106
2040044C	PWRK\$KNBDAEMON	HIB	12	1812	0 00:00:04.93	341	393
2040044E	PWRK\$LICENSE_R	HIB	11	244	0 00:00:04.31	538	479
20400454	PWRK\$LMMCP	HIB	11	2058	0 00:00:04.85	1247	505
20400458	PWRK\$LMSRV	HIB	11	15723	0 00:00:09.57	3008	2234
20400452	PWRK\$MASTER	HIB	6	186	0 00:00:04.04	460	175

Should one or more processes be missing, issue the following command:

```
$ ADMIN /ANALYZE [ /SINCE[=DATE-TIME] ]
```

This will give you an error message and the name of the failing process, if any.

You may also find license related errors here.

### Example 1

`$ PWSHOW` tells you that all the `PWRK$LM` processes are missing and `$ ADMIN /ANAL` tells you that the process `PWRK$LMSRV` was the one that exited first.

Then the next step would be to take a look at the process log file:

```
$ TYPE PWRK$LLOGS:PWRK$LMSRV_<NODENAME>.LOG
```

Here you can find the signature of the real problem. Please report this to your support centre to get a solution.

## Example 2

Let's look at another scenario, same report from a user; server is unresponsive.

\$ PWSHOW tells you that all processes are there. So there is no failure, its just slow.

The next step would be to take a better look at the output of \$ PWSHOW, and maybe repeat the command.

Look at the process PWRK\$LMSRV, this is the process that is handling all File and Print Server traffic.

Is the PWRK\$LMSRV process in HIB state, in LEF state or in RWSCS state?

1. If the state is HIB(ernating), take a look at \$ PWMON CLIENT /TOP=OPERATIONS. It is likely that you will find a PC client here who is heavily communicating with the server. You should check what the PC is doing. For example, it could have turned on virus scanning of network drives.
2. If the state of the PWRK\$LMSRV process is LEF, you may have a problem with one of your disks. This state means that it is waiting for an I/O completion. Start looking at \$ SHOW DEVICE D. This may indicate, for example, a device in trouble like mount verification. Another approach to disk related issues could be \$ ANAL /SYSTEM. Then, within SDA>, do a "SHOW PROCESS PWRK\$LMSRV /CHANNEL" and look for busy channels to a disk or file.
3. If the state is RWSCS it indicates the process is communicating to another node in the cluster. Very often this is locking, you can take a look at \$MONITOR DLOCK and \$MONITOR RLOCK. If \$MONITOR DLOCK reports high activity and \$MONITOR RLOCK shows little lock remastering, there is nothing you can do about it, the server is working hard. If MONITOR RLOCK shows high lock remastering, you could contact your support center to discuss a change of system parameters to influence lock remastering. What is considered high lock remastering depends completely on your system capabilities, specially the channel between the systems.

## Example 3

Your user cannot get a connection to his shares but others can still work.

This could have various sources.

### Example 3a

\$ ADMIN /ANALYZE /SINCE can give you a lead to a licensing issue. You could see an error like:

```
Event Time: 18-JUN-2003 20:35:45.96      Node:  UTURBO
Process Id: 20A00285
Event:      No server license for client - access denied
Event Source: LAN Manager Server
Event Class: Warning
```

```
Client: SMETSERSDENNI
```

This error report will also show the PC that is struck by the error which you can use to verify the user that is complaining.

In such a case, check with the command: \$ PWLIC to see if you have sufficient licenses available.

```
UTURBO > PWLIC
```

```
Advanced Server for OpenVMS (V7.3-120A): Server-Based License Report:
```

License	Total	Cluster Use	Available
PWLMXXXCA07.03	1750	29	1721

Note here, PWLIC displays cluster-wide license usage.

The available count could have reached 0.

If this is the case, your action depends on whether you have the license server running.

If you have, the process PWRK\$LICENSE\_S is running, then you may consider moving some licenses to the group called server-based.

If you do not use the license server, you will have to buy more licenses or reduce the number of clients that is using the system. Reducing the number of clients cannot be achieved by configuring the server for fewer clients since you cannot tell how many sessions each client will create to your server. You can only achieve this by changing the usage, for example, turn off browsing. Please be careful when you consider turning off browsing. A PDC will need the browser service but a member server will normally not need it. In case of a Backup Domain Controller it completely depends on your network topology.

To turn off browsing use the REGUTL utility to create or set the MAINTAINSERVERLIST key to NO.

### Example 3b

```
$ ADMIN /ANAL /SINCE does not show any license errors.
```

You may have configured your server for too few clients. To monitor this you have to check a few things.

First start with \$ ADMIN SHOW SESSIONS and check the current number of PCs using your system.

This can give a rough indication but does not include inter-domain sessions and browsing sessions.

Next thing to check is the protocol usage: \$ NBSHOW KNBSTATUS

The output of this command shows the sessions that are in use over TCP/IP ONLY.

There is a similar command for the NETBEUI protocol: \$ NBSHOW NBSTATUS

```
Take a look at the line: Sessions: In use: 99 of 100;
```

This will tell you if you ran out of session slots. If these numbers get close, like in the example above, you will have to reconfigure your server and increase number of clients.

To do this start \$ ADMIN /CONFIG, you will find the number of clients as "Client Capacity" on the first screen.

Please also note that the list of sessions at the end of the output of NBSHOW (K)NBSTATUS only lists maximal 112 sessions. The line containing "Sessions: In use:", mentioned above, is the one you should check.

You do not have to shutdown your server immediately to do this reconfiguration but you will have to reboot to effectuate it.

## **Documentation**

Documentation for Advanced Server is available on the internet. Please look at the HP documentation site; <http://h71000.www7.hp.com/doc/advserv73.html>

## **Futures**

There will be continuous development of Advanced Server for OpenVMS.

The first ECO release ( V7.3A-ECO3 ), which is planned for July 2004, will have a special cache around the SpoolSS printing interface to improve the performance of NT-style printing.

There is another eco release for version 7.3A planned to add support for OpenVMS Alpha 8.2.

The first major step will be to bring the current product on the new Itanium release of OpenVMS.

Another step that is planned is the addition of Active Directory integration. This will mean that Advanced Server can publish its resources in the Active Directory.

## **Ask The Wizard**

Should you have just a question about the product, not an error report, you can ask this to one of our wizards on: <http://h71000.www7.hp.com/wizard/>