FUJITSU

WHITE PAPER

FlexFrame™ for SAP®
Version 5.0A

THE FLEXFRAME INFRASTRUCTURE SOLUTION - TECHNICAL WHITE PAPER

INTRODUCTION

This document describes the technical aspects of the FlexFrame infrastructure solution as of version 5.0, which replaces the predecessor version 4.2B.

All technical aspects described in this document and this document itself is subject of change without further notice.

TARGET AUDIENCE

The document addresses consultants, administrators and platform decision makers.

We assume that the reader of this document has technical background knowledge in the areas of operating systems Linux, IP networking and SAP basis.

DOCUMENT HISTORY

Document Version	Date	Changes
2.0	2009-02-02	Update for FF 4.2
2.1	2009-03-11	Including customer-specific VMware-integration, adjustments for backup for Oracle with NetWorker
2.2	2009-06-10	Rebranding
2.3	2009-12-10	4.2B, RBAC, PCL → Linux-HA and adaption to DI-Extranet, RX300 S5, Elimination of Solaris
2.4	2010-03-15	Support SLES10 SP3
2.5	2010-04-22	Update for FF4S 5.0
2.6	2010-09-02	New SAP-Certificates
2.7	2010-10-11	Introducing Nexus Access Layer Switches

CONTENTS

1 The FlexFrame Infrastructure Solution

The FlexFrame™ infrastructure solution is a revolutionary approach to run complex SAP® infrastructure solutions with higher reliability and efficiency.

At the same time some major changes to the configuration paradigms for infrastructures have been implemented. These changes are:

- A shared operating system booted over IP networks for the SAP servers.
- Decoupling of application software and operating system, called virtualization of SAP software.
- Shared Network Attached Storage from NetApp or EMC. Additionally SAP database data may reside on either Network Attached Storage or SAN Attached Storage from either NetApp or EMC.
- FlexFrame Autonomous Agents (FA Agents) providing revolutionary mechanisms to implement high-availability functions without cluster software.

2 Components

The concept of FlexFrame for SAP (FF4S) consists of several components which implement state-of-the-art functionality. Together with new components, such as the FlexFrame Autonomous Agents, the whole solution is far more than just the sum of its components. A major part of the benefits consists in a dramatic reduction in day-to-day operating costs for SAP environments.

It is, of course, possible to use parts of the FlexFrame infrastructure solution in individually designed implementations[1]. In such customer-specific configurations, professional service and maintenance have to be considered separately.

2.1 Hardware

A typical FlexFrame environment consists of the following types of hardware:

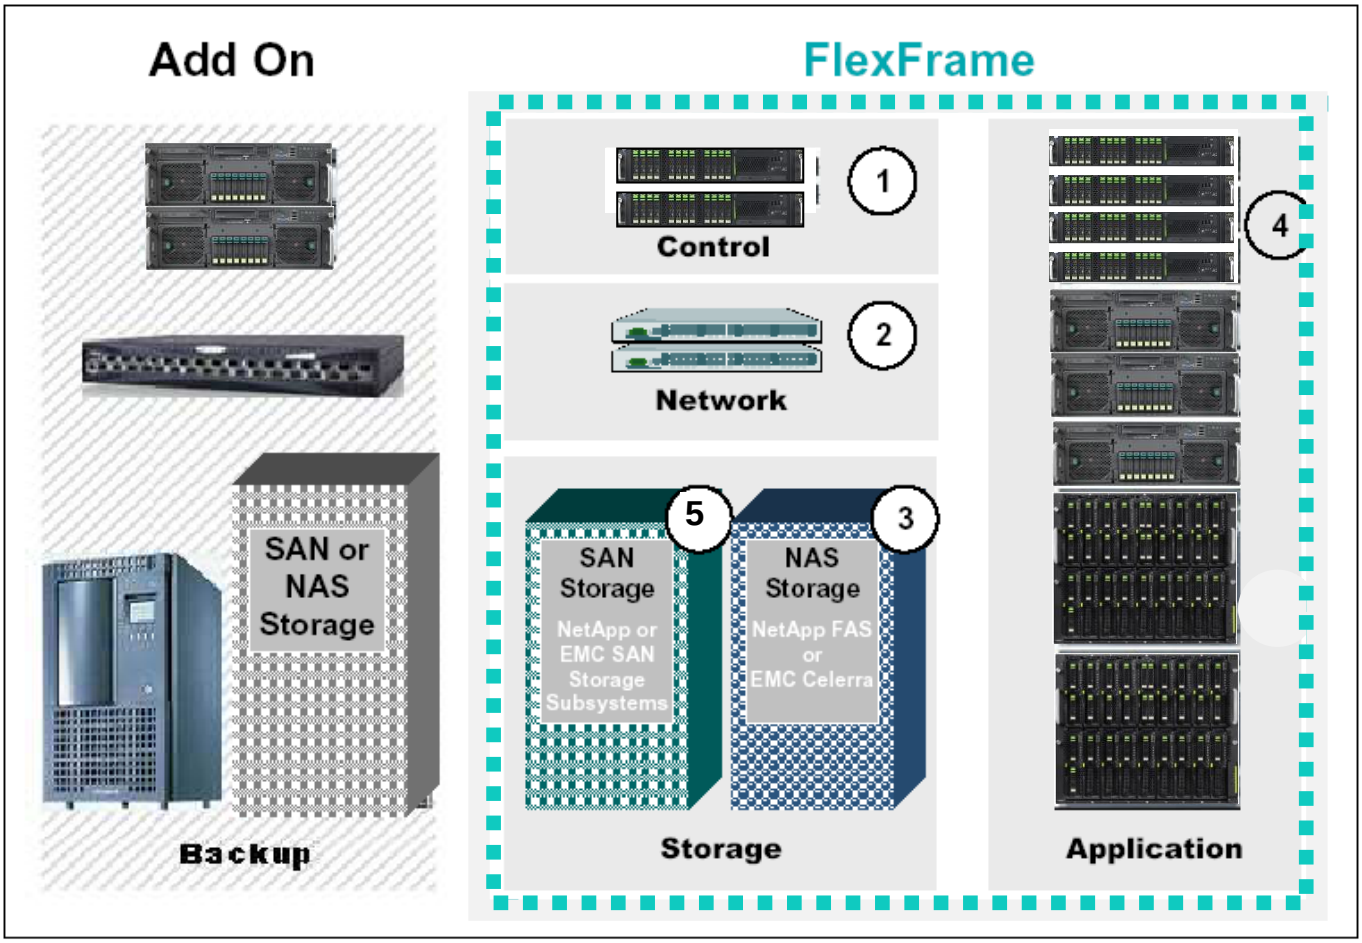

1. Two Control Nodes – the Control Center: PRIMERGY RX300 S6, RX300 S5, RX300 S4 or RX300 S3 with SuSE Linux Enterprise Server SLES10 64Bit installed on local disks.
2. At least two (equal) network switches Cisco® Catalyst 3750G or 3750E, optionally pairs of Cisco® Nexus 50x0 network switches in addition.

[1] Such implementations are classified into “standard +” Solution projects, where a standard special release request has to be made to the product owner, and “extended” (enhancements that later become mainline) and “customized” (enhancements without path into the mainline) Solution projects, where the project manager needs to involve the product owner in order to get a proper solution, to generate the mutual back-to-back agreement

3. Network Attached Storage with one NetApp FAS System or EMC Celerra™ hosting shared OS file systems and application data.
4. Intel- or AMD-based PRIMERGY Server (standard rack server or blade server) serving as Application Nodes with SuSE Linux Enterprise Server 10 or SuSE Linux Enterprise Server 11 (Shared OS).
5. Optionally SAN Storage Subsystems (from NetApp or EMC) containing LUNs carrying SAP Databases and Log Volumes.

For detailed information about the hardware supported in a FlexFrame environment, see the FlexFrame Configuration Guide and the current FlexFrame Support Matrix (https://partners.ts.fujitsu.com/com/products/infrastruc-solutions/FlexFrame).

Any other functions, such as backup, can be implemented separately as an add-on to FlexFrame and need dedicated hardware, operating system, high availability, professional service and support etc.

2.2 Software

The FlexFrame infrastructure solution consists of both hardware and software. To enable the complete environment to function properly, the entire software set is strictly defined. Anything other than the software components listed below is not part of FlexFrame. This also applies if software from the list below is missing, is installed in other versions than specified below, or if software other than the current SAP components is added.

2.3 Shared Operating System

One major aspect of FlexFrame is its shared operating system. Sharing in this case means that the very same files of essential parts of the underlying operating system are used to run multiple servers. This part of the file system is mounted read-only, so none of the Application Nodes that run the actual applications can modify it. Server-specific information is linked to a file system area that is server-specific and mounted read/write. The shared operating system (normally named “image”) is kept on a Network Attached Storage system from either NetApp or EMC.

No.	Hardware	OS	Software	Services
1	Control Nodes: 2 x PRIMERGY RX300 S6, or 2 x PRIMERGY RX300 S5, or 2 x PRIMERGY RX300 S4 or 2 x PRIMERGY RX300 S3	SLES10 SP3 (x86_64)	FA Agents (Control Agents) V9.0, FlexFrame 5.0A File System Image CN, ServerView etc.	TFTP, DHCP, LDAP, (SAPROUTER), etc
2	Network switches: n*m Cisco Catalyst 3750x 2, n≥1, m≥2 n*2 Cisco Nexus 50xx [2], n≥0	IOS (proprietary) NXOS (proprietary)	(as delivered)	

[2] allowed types according to FlexFrame Support Matrix

No.	Hardware	OS	Software	Services
3	Network Attached Storage: – one or more NetApp Filer heads (FASxxxx), disk shelves as required [3], hosting shared OS file systems and application data or – one or more EMC Celerra NSxx, disk shelves as required [3], hosting shared OS file systems and application data.	ONTAP [4] (proprietary) DART [4] (proprietary)	ONTAP [4] NetApp Tools DART [4] EMC Tools	NFS, SnapRestore optional Cluster components, FlexClone, SnapVault, SnapMirror NFS, …
4	SAN Storage Multipath SW		SLES10 / SLES11: – DM-MPIO integrated Multipath SW	HA Services
5	SAN Storage Volume Manager		SLES10 / SLES11: – LINUX Volume Manager LVM2	Volume Management Services
6	Intel- or AMD-based PRIMERGY server (standard rack or blade server)	SLES11 SP1 (x86_64) and / or SLES10 SP3 (x86_64)	FlexFrame V5.0A File System Image, FA Agents (Application Agents), SAP Applications, Database	SAP & DB Services

2.3.1 Shared OS Boot Concept

The figure below shows the boot process of a FlexFrame Application Node (PRIMERGY/Linux).

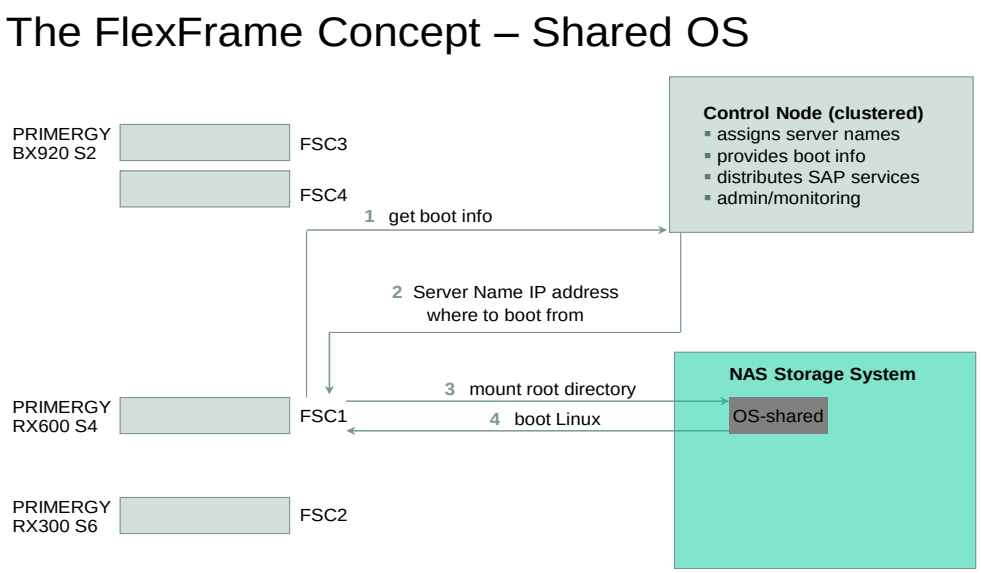

[3] The amount of disks required for customer-specific FlexFrame configurations can be determined together with Fujitsu Customer Support Filer Sizing Team
[4] The currently supported ONTAP resp. DART version is kept up-to-date in the FlexFrame for SAP Support Matrix

2.4 Control Nodes

Every FlexFrame environment always includes two Control Nodes. Their purpose is to be a single point of control for the Application Nodes, as well as to check and manage the proper operation of the Application Nodes.

The two Control Nodes are also called the Control Center.

Control Nodes do not run SAP software (with the exception of saprouter, as an option). As of FlexFrame for SAP 4.2B they exclusively run SuSE Linux Enterprise Server Version 10), installed on local disks. Control Nodes provide and run services such as:

- Timeserver for the complete FlexFrame environment
- FlexFrame Autonomous Control Agents
- Web server to provide the Control Agents' user interface
- Assignment of IP addresses to Application Nodes using DHCP (PRIMERGY server)
- TFTP server for the boot process
- LDAP server
- ServerView Operations Manager for monitoring and analyzing PRIMERGY servers
- saprouter (optional, limited to SAP OSS connection or administrative purposes)
- Preparation for Squid Proxy and DNS server for communication between clients and SAP services
- StorMan for the execution of dynamical LUN masking via SMI-S interfaces

Control Nodes are of the server type PRIMERGY RX300 Sn, n = 6, 5, 4 or 3.

2.5 Application Nodes

Application Nodes are the workhorses of the FlexFrame infrastructure solution. The Application Nodes offer CPU and memory and run database and SAP services. They need local disks, though they are only used for swap space and do not contain any other data.

For FlexFrame Version 5.0A the principal types of Application Nodes are:

- PRIMERGY server running SLES10 (64bit)
- PRIMERGY server running SLES11 (64bit)

Admissible servers have to be approved for SAP by Fujitsu.

During the boot process using PXE®, each Application Node is identified using the hardware address of its boot interface
(MAC address). The Control Node assigns an IP address to it and supplies the operating system via the network.

Selected file systems (especially the root file system "/") are mounted over the network read-only.

If, for any reason, an Application Node needs to be replaced or added, only a handful of settings need be adjusted to integrate it into the FlexFrame environment.

Intel's® PXE technology is implemented in Fujitsu's PRIMERGY servers and allows booting over the network. DHCP is used with static MAC address relationship for all the Application Nodes.

2.6 VMware® ESX server integration into FlexFrame for SAP

Virtualization concepts play a fundamental role for the customer value of the dynamic infrastructure FlexFrame for SAP. Besides the concept of virtualization of the SAP software by decoupling application software and operating system and the concept of virtual IP addresses enabling the switch of services from one host to another, FlexFrame now offers the use of virtualized servers as Applications Nodes.

This milestone firstly had successfully been reached on a project basis for some FlexFrame for SAP customers. The used technique is an integration of VMware® ESX server (resp. the VMware® ESXi server) into the customer's FlexFrame for SAP. This integration now has become a turnkey part of the solution FlexFrame for SAP with its version 5.0.

Starting Point is the general virtualization design:

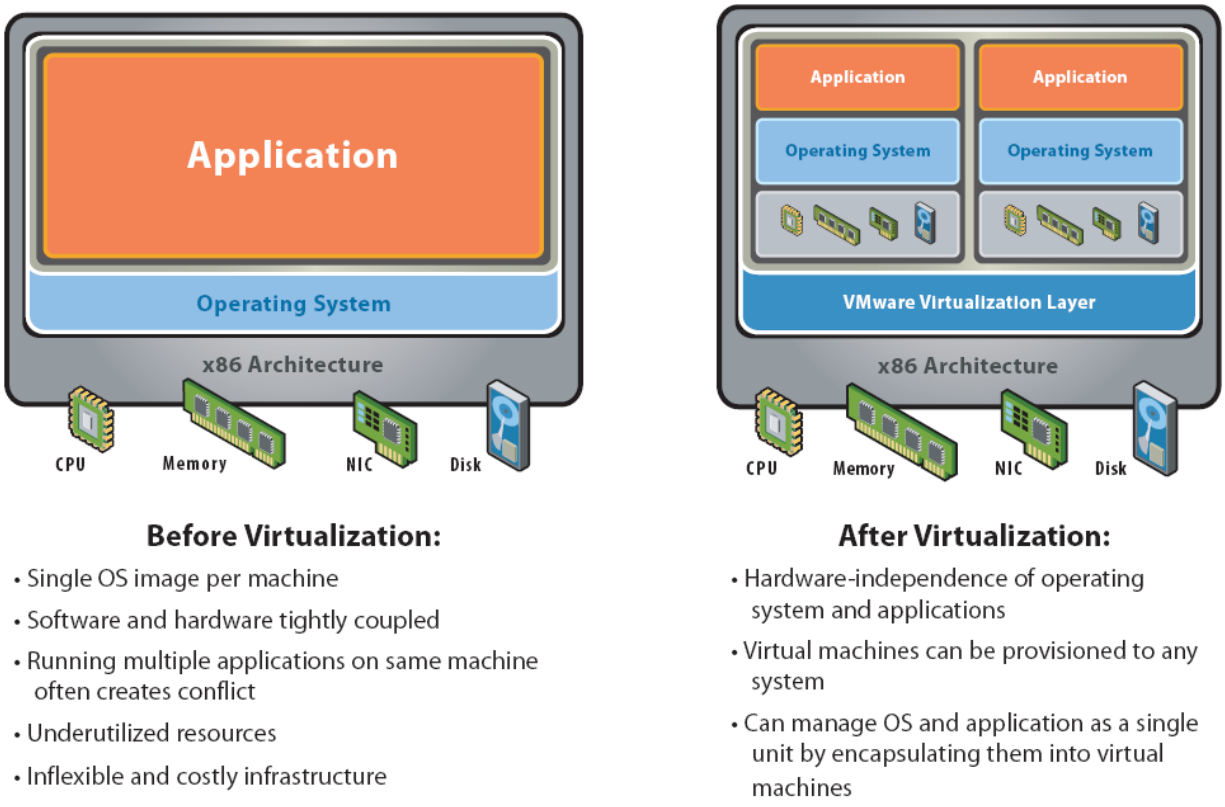

Before Virtualization:

- Single OS image per machine
- Software and hardware tightly coupled
- Running multiple applications on same machine often creates conflict
- Underutilized resources
- Inflexible and costly infrastructure

After Virtualization:

- Hardware-independence of operating system and applications
- Virtual machines can be provisioned to any system
- Can manage OS and application as a single unit by encapsulating them into virtual machines

Allowing Application Nodes with and without VMware resp. ESX server the obtained picture in FlexFrame for SAP is

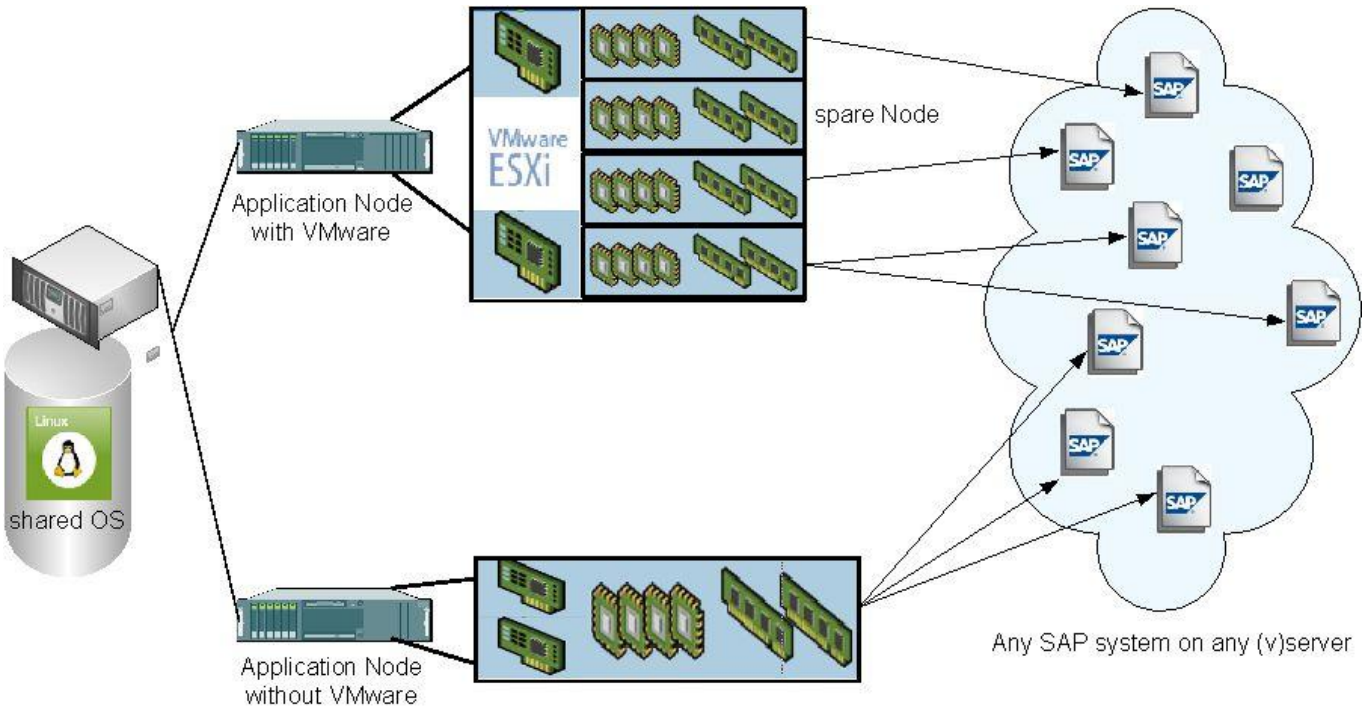

FlexFrame for SAP makes it possible for an SAP application to be started on any server in the FlexFrame pool. This can be a physical server or a virtual server under VMware. An SAP application can be relocated from a virtual server to a physical server and vice versa at any time. As a result, the application can be moved to the system with the best suited resources without any additional installation work. For example, several applications with low loads are run in parallel on several virtual servers, although if an application's load is high, a complete physical server is assigned to it. If this is not sufficient, a server with more or more powerful processors and main memory is chosen.

All in all, the available server capacities can thus be utilized better and thus reduced. Unused resources can also be switched off. Both measures cut procurement and power costs. Flexible allocation is also a great aid if a new SAP application does not run correctly on a virtual server. It is then moved to a physical server to test whether the problem is the result of virtualization or due to other reasons.

Management

All installed and configured resources are controlled from a central instance, the FlexFrame Control Center. It can essentially be assumed that an SAP environment consists of a mix of virtual and physical servers. Uniform management of all resources – virtual or physical – is therefore an important requirement for the IT infrastructure. Central management of the SAP services also offers many advantages in the event of a fault. If a server has failed, the SAP application is automatically booted on a substitute system. SAP applications are also restarted on virtual servers. The basis for this is a FlexFrame agent that is started on all application nodes when the operating system is booted and reports to the Control Center that it has been started and is running successfully. This agent allows SAP-specific problems to be detected and countermeasures to be initiated. As a result, not only availability of a virtual or physical server is monitored, but also whether the application is running smoothly and its utilization. That means administrators have an overview of all server resources and applications at any time and can also change the assignment of resources to applications. There is thus end-to-end control of the entire system.

2.7 Role Based Access Control (RBAC)

In computer systems security, role-based access control (RBAC) is an approach to restricting system access to authorized users. Within an organization, roles are created for various job functions. The permissions to perform certain operations are assigned to specific roles. Members of staff (or other system users) are assigned particular roles, and through those role assignments acquire the permissions to perform particular system functions.

RBAC is an add on option to FlexFrame for SAP and needs to be licensed separately for usage.

In order to configure and administrate a FlexFrame environment numerous programs have to be used, nearly all of them needing root permissions to do their job. So the previous FlexFrame versions for example did not allow to separate administration for different pools (and often a pool corresponds to a customer in outsourcing environments)!

With RBAC a FlexFrame administrator

- need not have “root” privilege any more
- can be a named account on the Control Center
- can be restricted for different administration scopes (pool, sid, ..)
- will be logged with all its interface actions done by FlexFrame programs

RBAC uses two data structures to define access permissions:

- roles
  A role contains all permitted FlexFrame programs and operation modes for this role.
- user roles
  A user role maps FlexFrame administrator roles by pool to a user role name. This name has to match to an account or group name. If the account name matches any user role no more user roles are scanned. If no user role name matches the account name all the groups of the account are used to scan for user roles. The permissions of the group name matching user roles are added to a composed virtual user role.

2.7.1 RBAC Notices

- RBAC does not do any system hardening on the FlexFrame Control Center but can be used to limit the FlexFrame administration to named accounts without direct “root” rights
- RBAC is not able nor designed to substitute sudo. It will only grant permissions to scripts/programs living at /opt/FlexFrame/bin. These scripts / programs have to use the RBAC API to be able to participate.
- RBAC does not cover the normal system administration or some expert FlexFrame interventions which need root privileges.
- Therefore be aware which person or group you define to be a FlexFrame administrator – regardless using RBAC or not. Any not authorized person may be a security gap in your organization. This will not be closed by RBAC.

2.7.2 RBAC standard role definition

Out of the box FlexFrame comes with six role definitions. Three of them are for configuration/administration and the other three for observation. They are divided into SID specific, pool specific and main FlexFrame and do not overlapping in most cases. This means to have grants to SID and pool specific tools both roles have to be added to user role. To get a non root user with root like FlexFrame permissions all three configuration/administration roles have to be added to the user role of the user.

The three configuration/administration roles are:

- ffadmin
- pooladmin
- sidadmin

The three observer roles are:

- ffobserver
- poolobserver
- sidobserver

The intention of the predefined rules is to give a non root user the ability to do all the SID or pool specific FlexFrame jobs or to do FlexFrame environment related work. These are divided into a “read only” observer part and an administrative part.

2.8 Use of SAN Attached Storage in FlexFrame Environments

2.8.1 SAN Hardware Configurations

The integration of SAN Storage Subsystems in FlexFrame allows benefiting from the performance advantages of fibre channel connections compared to 1Gbit/s Ethernet connections and thus (as an alternative to 10Gbit/s Ethernet connections) closes the possible performance gap in present-day NAS environments approaching some FlexFrame customers in the near future.

The support of SAN attached storage within FlexFrame further enables building up FlexFrame infrastructure solutions in existing customer environments using SAN as the standard storage technique for a long time.

While data areas like operating systems, application software and commonly used data of FlexFrame infrastructure solutions have comparatively low need for data throughput, these areas remain on a network attached storage (NAS) and thus still benefit from the flexibility of the present solution, the central data areas of type database and loggings (sapdata and saplog) needing high data throughput may be shifted to SAN storage.

The next figure gives a simple overview of a FlexFrame SAN configuration:

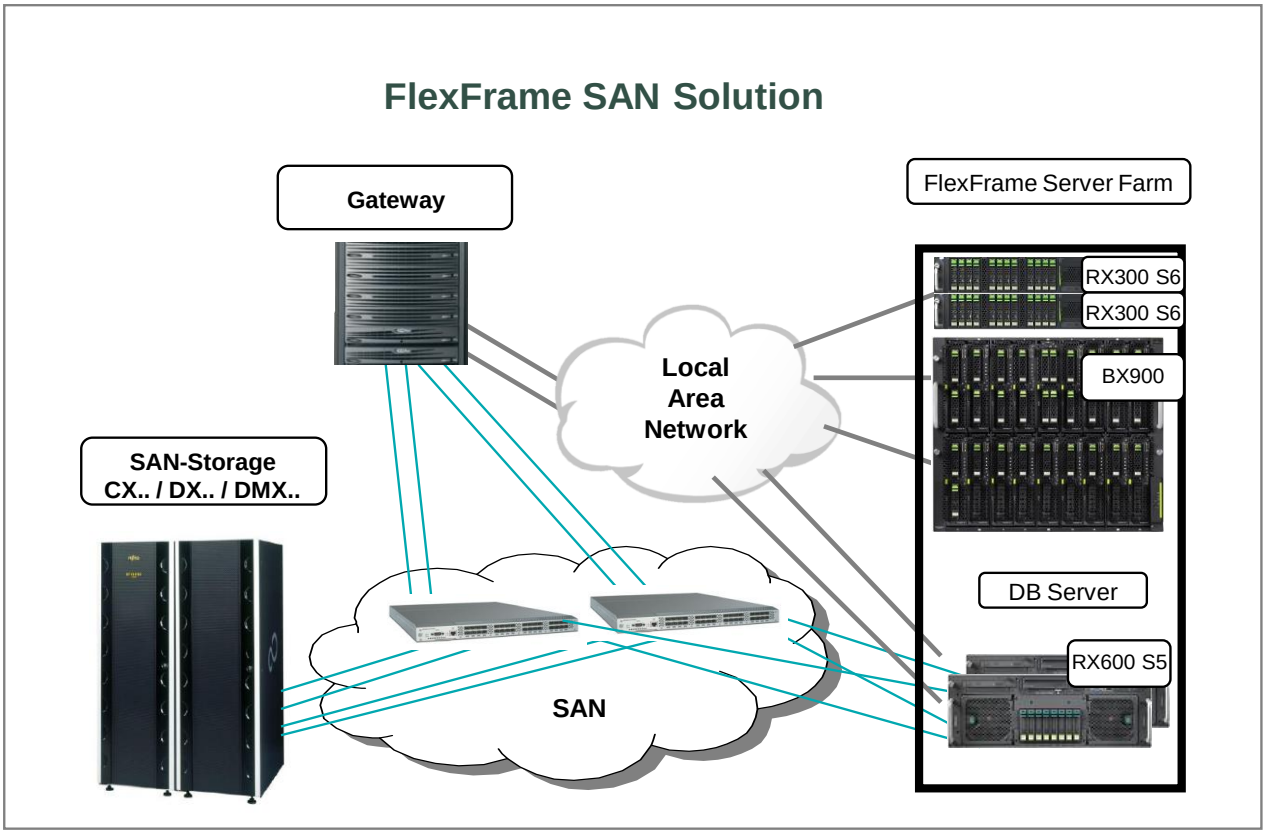

While the FlexFrame NAS attachment is supported by the automatic configuration of FlexFrame access layer switches (including their wiring) the corresponding configuration and wiring of SAN switches concerning the FlexFrame SAN attachment is not supported by automatic configuration within FlexFrame. This part has to be done manually according to the customers storage requirements and is independent of FlexFrame.

2.8.2 SAN Software Configurations

In order to be able to use SAN configurations in the FlexFrame environment a few SAN management functions are supported and have to be integrated in FlexFrame. For high availability and high performance reasons **SAN multipathing** is supported in many forms. For logical volume management tasks such as allocating space on storage systems, concatenate, stripe together or otherwise combine partitions into larger virtual ones, apply RAID

techniques and others most customers deploy a **Volume Manager** software, and in case of FlexFrame SAN solutions also a standard set of Volume Managers has been qualified.

The following table summarizes the basic facts about the FlexFrame 5.0A supported software tools of type

- SAN Storage Multipath SW
- SAN Storage Volume Manager
- SAN Storage Administration and Management Services

Product / Component	Type of Storage SW	Version	Provider	Part of FF4S delivery	Installed on
DM-MPIO SLES10	Multipath	Part of SLES10	Novell	yes	Linux-AN
DM-MPIO SLES11	Multipath	Part of SLES11	Novell	yes	Linux-AN
LVM2	Volume Manager	Part of SLES10 / SLES11	Novell	yes	Linux-AN
StorMan	Storage Provisioning	2.1	Fujitsu	yes	CN
SMI-S Provider EMC	Provider for SMI-S V1.1	4.1	EMC	no	CN or special storage administration server

The use of these software tools is optional within FlexFrame. Other tools of these types have not been tested and qualified for FlexFrame for SAP. The tools which are not delivered with FlexFrame can be ordered by the corresponding provider and have to be installed by the customer.

2.9 Boundary of a FlexFrame Environment within a Data Center

The term boundary is meant here to name the borders of FlexFrame in the sense of distinguishing which data center hardware belongs to a FlexFrame landscape and which does not. In customer configurations mostly there are well defined boundaries of a FlexFrame environment within the data center. Any Application Node belongs completely to FlexFrame and is reachable from the outside only via the FlexFrame Client LAN (one of the 4 virtual LAN segments of FlexFrame, see section VLAN Segments). The NAS storage subsystems often likewise belong to FlexFrame completely - in some configurations the NAS systems are subdivided into a part (diskshelves, disks, controller, ports, etc) belonging to FlexFrame and a part belonging to the rest. Also SAN attached storage might be used within and without FlexFrame as well; then the LUNs and/or ports are separated at least by Fibre Channel SAN zones or by LUN Mapping / Masking. The boundary of FlexFrame to the outside world essentially comprises of the uplinks of the LAN switches of the FlexFrame Switch Groups.

Physically a FlexFrame environment is concentrated in a set of PrimeCenter racks, which contain assemblies of hardware systems normally belonging completely to FlexFrame, although sometimes due to shortage of space non FlexFrame rack elements are built in FlexFrame racks.

With FlexFrame version 4.2 a new "switch blade" type is usable for BX600 blade servers: the Gbit Ethernet Pass-Thru Blade 10/10. This Pass-Thru Blade enables the dedicated Gbit Ethernet access to each single blade server and allows to have FlexFrame-dedicated and non-FlexFrame-dedicated blade servers in one blade chassis. This feature "shared utilization of BX600 for FF4S and other applications" is only available by a request for special release

A corresponding feature “shared utilization of BX900 for FF4S and other applications” is no integrated part of FlexFrame V5.0A but could be realized by a customer specific project solution through the additive use of a second pair of switch blades (in a second BX900 fabric).

3 Central Administration via LDAP Configuration Database

LDAP is used as the central information service for all shared OS nodes within a FlexFrame environment. The Control Nodes are used as LDAP servers. The LDAP database is located on shared file systems mounted from the NAS storage. The Application Nodes are configured as LDAP clients. LDAP requests from Application Nodes are restricted to the data of their own pool.

LDAP provides host-related network information such as:

- net boot
- automount
- user authentication
- groups
- host names and IP addresses
- shared services
- networks and netmasks
- LDAP client profiles

Additional information about configuration data is only applicable for Control Nodes. These are used FlexFrame internally to add, remove or modify the configuration of Application Nodes or SAP services.

As of FlexFrame for SAP 4.2B the standard high availability and failover solution “Linux-HA cluster” is used for the two Control Nodes (replacing the former PRIMECLUSTER solution). This feature also guarantees, that after breakdown of a Control Node the cluster provides its IP address and the corresponding LDAP resource (LDAP-master or LDAP-replica) to the other node and makes it available again for the entire system.

4 IP Network

The network is the backbone of the FlexFrame infrastructure solution. Communication between the various nodes and storage is done exclusively over the IP network infrastructure. This is used both for communication between server(s) and client(s) and for delivering IO data blocks from the NAS (Network Attached Storage) to the server.

The IP network infrastructure is essential for every FlexFrame configuration. FlexFrame is designed with a dedicated network for connections between servers and storage. This network is reserved for FlexFrame internal traffic only. By default the Client LAN (see section VLAN Segments below), is routed outside FlexFrame network for connection to the existing corporate network. Also other LAN segments may be routed to outside networks.

The FlexFrame internal physical network configuration is redundant to ensure reliability.

The network bandwidth has to be at least 1 Gbit/s for all components to handle intense data traffic. Selective network connections (for instance the connections from a FlexFrame NetApp NAS system or from a set of 10GbE-capable rack servers or blade servers (via switch blades) to the FlexFrame access layer switches) maybe chosen as 10Gbit/s connections (new feature in 4.2).

A number of virtual network segments are configured on the basis of this physical network.

4.1 Network High Availability – LAN Failover

The term "LAN failover" describes the ability of the FlexFrame environment to use a logical network interface that consists of several physical network interface cards (NICs) or ports, which in turn use redundant network paths (cables and switches). When a network component (port, NIC, cable, switch etc.) fails, the network management logic switches over to another path via a second port or network interface card. If only one NIC with two ports is present, and the NIC fails (on both ports), then a failover to another server will be initiated.

4.2 VLAN Segments

With FlexFrame a virtual network concept is established for providing higher availability as well as increased flexibility for the FlexFrame environment.

This concept relies on virtual LAN (VLAN) technology that allows multiple virtual networks to be run across a single network wiring system. Additionally, in order to ensure high network availability, LAN bonding is used on every node. This includes a double switch and wiring infrastructure to keep the whole environment working, even when a network switch, cable or NIC fails. Furthermore, the network segments are prioritized to ensure that important connections are given preference.

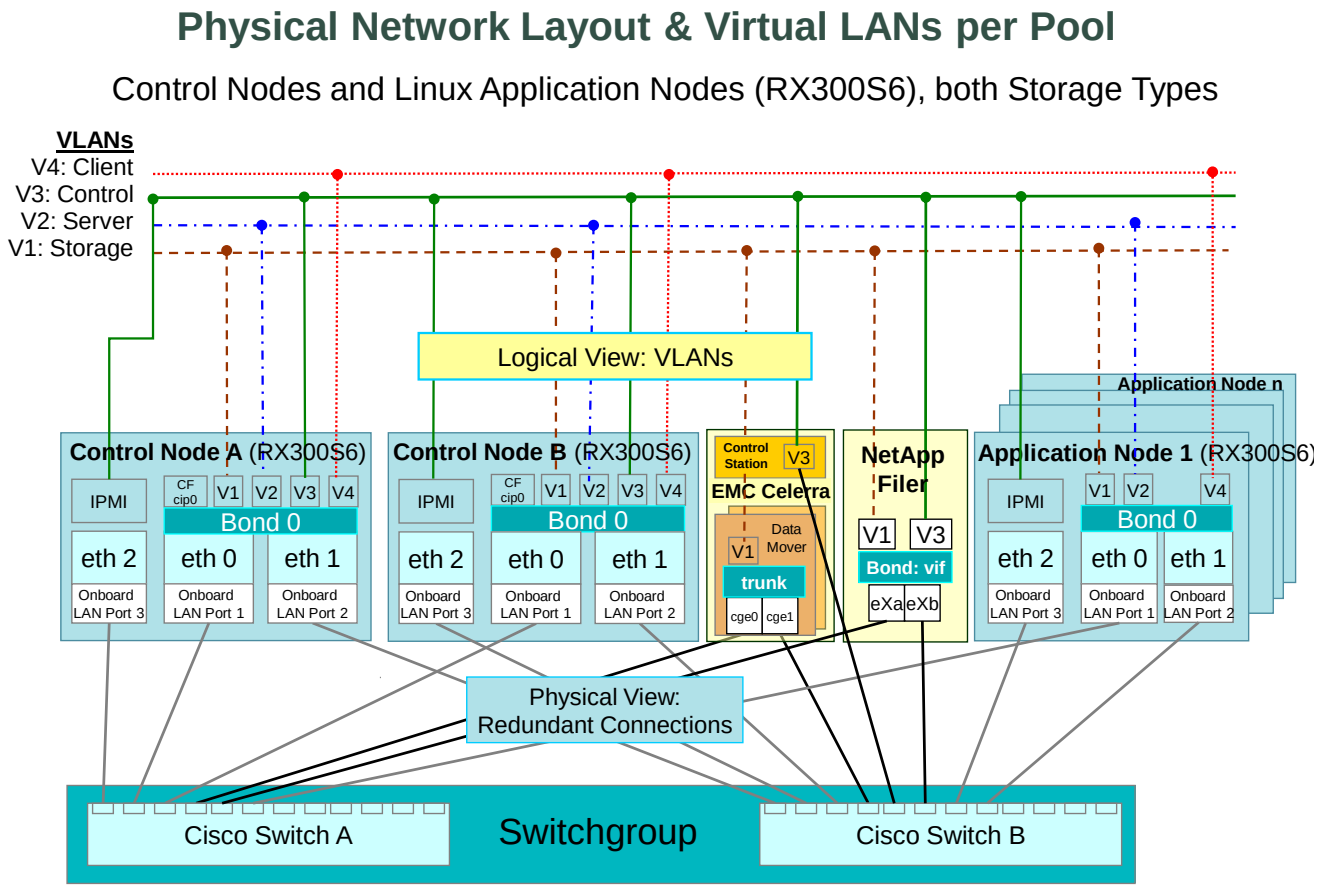

Similar to former FlexFrame versions, typically four networks are used within one FlexFrame system or within a single FlexFrame pool. FlexFrame uses virtual network segments. These virtual Networks run over the same physical Network. The physical network is built with redundant NICs, using bonding on Linux systems.

The figure outlines the basic network segments of a typical FlexFrame environment with Linux-based Application Nodes.

The following virtual network segments are mandatory for FlexFrame:

- **Client LAN (one for each pool)**
  The purpose of the client LAN segment is to provide dedicated user connectivity to the SAP instances. This segment also allows administrators to access the Control Nodes.

- **Control LAN (one per FlexFrame installation)**
  The control LAN segment carries all administrative communication for the IPMI interfaces (intelligent

platform management interface) or similar interfaces.
Note: Control LAN access to the Application Nodes is required for power on/off/cycle interventions.

- **Server LAN (one for each pool)**
  The server LAN segment is used for the SAP instances to communicate with each other and with the databases instance.

- **Storage LAN (one for each pool)**
  The storage LAN segment is dedicated for NFS mounted file systems. Files of type Application Node operating system, application software and commonly used data of FlexFrame infrastructure solutions are located on the centralized NAS storage. Database files, files from SAP instances, database instances may also reside on the centralized NAS storage or are located on SAN attached LUNs (and in that case do not stress the Storage LAN).
  The storage LAN is also used for netboot of Application Nodes.

4.3 Network Switches

Network switching components play a very important role within FlexFrame. Currently the following switches are tested and supported as FlexFrame access layer switch:

- Cisco Catalyst 3750G & 3750E
- Cisco Nexus 5010 & 5020

These switch families support VLAN technology for flexible configuration of the various network segments and QoS (Quality of Service) to prioritize traffic (e.g. higher priority for production systems compared to test systems etc.). For Catalyst models the Cisco StackWise® Technology is able to use cross-switch link-aggregation (this means the usage of ports of different switches within one stack). This concept is currently used for NAS system connections, switch group interconnects and switch blade connections. The Nexus counterpart of a Catalyst switch stack is the vPC domain consisting of exactly two Nexus switches.

As of FlexFrame Version 4.2 there are two different network speeds supported:

- 1Gbit/s
- 10Gbit/s

With the 3750G models you can only use 1Gbit/s network speed. With the 3750E models you can use 1Gbit/s and 10Gbit/s network speed. The Nexus models provide numerous 10Gbit/s ports and no others. The FlexFrame network architecture requires at least two 1GbE switches of type Catalyst, whereas additional 10GbE switches of type Nexus are optional.

For 1Gbit/s you can use copper or fiber ports. For 10Gbit/s you can only use fiber ports.

The 10Gbit/s network can be used for the connection of 10GbE-capable NetApp FAS Systems as well as for the connection of 10GbE-capable rack servers and blade servers and as uplink to the customer core network.

In FlexFrame network switches are grouped. Ports of an endsystem building a redundant connection are connected to ports of different members of the same switch group. A switch group consists of at least two switches of the Catalyst 3750G or 3750E switch family or of exactly two Nexus Switches.

5 Central Network Attached Storage

The storage for all Application Nodes is consolidated on one (or more) central Network Attached Storage systems with a connection based on the network file system protocol (NFS). The central NAS storage can be a FAS System from NetApp or a Celerra Network Server from EMC.

Both the NetApp and EMC implementations of the NFS (Networked File System) allow the same data files to be shared by multiple hosts and thus provide a built-in cluster file system.

The FlexFrame concept reduces the amount of "lost" disk space, since multiple SAP systems can optionally share the same set of disks. As the data grows, it is easy to add additional disks and enlarge the volumes without any downtime.

The storage configuration for FlexFrame concerning capacity and IO requirements are determined by the Fujitsu SAP Presales process.

5.1 Central Network Attached Storage on NetApp FAS Systems

The operating system of the FAS System is called *ONTAP*. The disks are grouped into RAID groups. Combining RAID groups create an aggregate that can be split into volumes of flexible size. A volume contains a file system (WAFL™ - Write Anywhere File Layout™) and can serve as an NFS (for UNIX systems) or CIFS (for Windows®) volume or mount points. The FAS System has NVRAM (Non-Volatile RAM) that buffers committed IO blocks. The contents of the NVRAM remain intact if the FAS System should suffer a power outage. Data is flushed to the disks once power is back online.

The minimum FlexFrame environment has at least the following volumes:

- vol0 (ONTAP, configuration of FAS System)
- volFF (one common volume and optionally several pool specific and/or SID specific volumes containing the shared operating systems, file systems of Application Nodes and other software)
- sapdata (database files)
- saplog (database log files)

The amount of disks required for customer-specific FlexFrame configurations can be determined together with Fujitsu's Global Competence Centre Architecture Services Team.

Usage of multiple Filer as well as Filer cluster is possible for redundancy and reliability. A FAS System can be clustered to protect data against the failure of a single FAS System by specific NetApp tools. This switching from one FAS System to its cluster counterpart is transparent to the FlexFrame Application Nodes.

5.1.1 Snapshots

When a snapshot is taken, no data blocks are copied; just the information where the data blocks are located is saved. If a data block is modified, it is written to a new location, while the content of the original data block is preserved. A snapshot can therefore be created very quickly since only a small amount of data need be copied. The use of snapshots does not decrease the throughput and performance of the storage system.

Snapshot functionality allows the administrator to create up to 255 views of a volume.

The most important use of snapshots is for backup purposes. For information about how snapshots are integrated in the FlexFrame Backup Solution for Oracle® see chapter FlexFrame for SAP Backup Solution for Oracle with NetWorker on page 28.

5.2 Central Network Attached Storage on EMC Celerra Storage Subsystems (NS Series)

The NS Series merges the power of EMC's storage platforms with the maturity of the DART (Data Access Real Time) operating system to provide best availability, flexibility, and manageability. The NS Series is available in both gateway and integrated models and support single, dual and quad Data Mover configurations. Upgrading from a single to a dual configuration is done online and non-disruptively. Dual Data Mover configurations can be deployed in Primary/Primary mode for performance-oriented environments or Primary/Standby for additional hardware availability protection. The four Data Mover models feature N+1 availability configurations. In FlexFrame environments an N+1 availability configuration with $N \geq 1$ is assumed (at least 2 Data Movers with at least one Standby).

The NS Series/Gateway allows customers to leverage their existing EMC storage investments for their IP Storage (NAS and iSCSI) requirements. All of EMC's gateway products support connection to the full line of CLARiiON / FibreCAT CX Series fibre channel arrays as well as Symmetrix models (DMX Series and predecessors).

6 High Availability with FlexFrame Autonomous Agents

The following sections describe how SAP and database services are made highly available using FlexFrame Autonomous Agents. The FlexFrame Autonomous Agents consist mainly of two different types of component, the control agents and the application agents. These components together monitor hardware and software availability, and can intervene if any component of the FlexFrame environment should not react properly.

This intervention feature is a key feature of the FlexFrame infrastructure solution. Seen from a client viewpoint, it provides automatic repair of almost any imaginable failure.

The main intervention mechanisms used are

- Stop and restart services
- Reboot node and restart services
- Switch over services to another node and shut down/reboot the former host
- Power-off a node (in case of damage)
- Displace services of lower priority for services of higher priority if needed
- Start a service on a node in addition to another service running there

6.1 Control Agents

The control agents run on the Control Nodes and basically monitor the application agents. If an application agent is no longer responding, the control agent has to react and restart the application agent's services on a spare Application Node. To avoid the split brain error, the control agent initiates a power-off sequence for the Application Nodes that failed. This power-off sequence is issued via system-specific power-off hardware and software facilities.

6.2 Application Agents

While an Application Node is booting, the autonomous application agents are started. Once an application (SAP instance, database etc.) is started, the application agents will start monitoring it. No extra configuration is required. If an application should fail or cannot be started, the agent will try to re-start it, and if this fails, the entire service is started on a different Application Node (spare node). Agents on spare nodes compete for such failed services and the "winner" starts the service. Failed services can also be started on active Application Nodes, for this purpose there exist sophisticated possibilities of add rules, replace rules and substitution rules matching customer specific circumstances.

6.3 Spare Application Nodes

Legacy concepts (non-FlexFrame) have dedicated clusters for fixed services. In those concepts a complete set of redundant failover hardware is required, which is idle as long as no error occurs.

In FlexFrame any number of spare nodes will stand-in if any Application Node fails.

In FlexFrame the concept of server pools and groups is new invented. In a server group, only such types of server are configured which can stand-in for each other. Depending on the number of SAP services and their requested SAPS power, a number of servers may run in a server group (see also chapter Pools and Groups on page 26.

Spare servers and servers running services of low priority are the failback resources to enable any SAP or database service that may fail to move to another node in the same server group.

Spare server could be individually allocated to a pool and could stand-in for any failure of an Application Node (with same OS image) in that pool. Additionally also pool-independent spare servers are possible, which are assembled in a special spare Pool named *Adminpool*, and which can stand-in for any failure of an Application Node (with same OS image) in any pool.

6.4 Control Node Availability

Control Nodes monitor and control the whole FlexFrame environment. Therefore Control Node functionality itself has to be highly available. The standard high availability and failover solution “Linux-HA” is used for the two Control Nodes. As of FlexFrame for SAP version 4.2B this Linux-HA cluster replaces the formerly used PRIMECLUSTER™ (PCL). Linux-HA is an integrated part of SuSE Linux SLES10.

6.5 Generic FA Agent – Autonomy and High Availability for User Specific Services

In many user scenarios additional SAP or non-SAP services shall be utilized on the application servers in the FlexFrame environment, so there is a need to integrate these services into the FlexFrame autonomy rules.

With the FlexFrame Autonomous Agents from FlexFrame for SAP it is possible to integrate other services into the autonomy scenarios. The functionality allows pool specific definition of autonomy rules for additional services.

In this case the rules for detection, reaction and high availability have to be defined in rule files. The new generic FA Agent functionality allows an adaptation to project specific requirements. For this purpose a generic service is defined through a set of parameters which are used for its identification and which generate the service states. The description and definition of a service is arranged in state, detection and reaction models.

7 FlexFrame Disaster Tolerant Configurations

The storage subsystems – NAS and SAN – used in a FlexFrame landscape offer a rich set of features and options to implement high availability and disaster tolerance solutions. While some features are inherent to the respective storage subsystems and offer a high level of data integrity, reliability and availability, others can be configured explicitly or may be optional features that must be purchased and licensed separately. The best practices recommended by the storage system vendors should be reviewed and applied according to the business requirements of the organization where the FlexFrame landscape is deployed.

When planning a disaster recovery solution for a FlexFrame landscape, an adequate solution must be selected according to the importance of the affected business and the related costs. To achieve the highest level of data currency on the alternate site, a solution using real-time host-based or storage subsystem-based mirroring should be implemented. A solution for the NAS storage system is always needed, as the FlexFrame landscape holds some important data on NAS. If some SAP systems are configured for SAN usage, a solution for the data residing on SAN LUNs is also needed.

Host-based mirroring is an ideal solution for a cross-site FlexFrame SAN configuration, as it ensures that an actual copy of the data is available on both sites and that data is accessed through a virtualization layer that makes it transparent to the application which of the copies is really used. This solution is actually independent of the vendor of the storage systems used in the FlexFrame landscape.

The recommended disaster recovery configuration for a FlexFrame landscape with NetApp Filers is a switched MetroCluster with SyncMirror. This implies that the two NAS systems from site A and B constitute in fact a single clustered NAS system with two clustered controllers and redundant disks that are spread across both sites. When NetApp SAN storage is also involved, a disaster tolerance for this section can be built with the mentioned host-based mirroring technique or alternatively with the also mentioned MetroCluster with SyncMirror.

Irrespective of the selected solution, a detailed operations guide must be created specifically for each customer and his configuration and requirements. And all of the discussed solution architectures have in common that they behave different in different failure or disaster situations – there are automatic failover reactions as well as situations where manual intervention is required for failover, and the failback is a manual procedure in any case.

And moreover if desired there exists the possibility to suppress automatic reactions completely.

8 Additional Management Functions of myAMC / Autonomous Agents

The following section describes additional functions of management options which help to fulfill requirements such as

- FlexFrame event and fault management expert
- FlexFrame performance and capacity management
- FlexFrame accounting management
- FlexFrame reporting option

These functions are generally useable within FlexFrame for my SAP. They are based on the raw information which is provided by the FlexFrame autonomous agents.

For detailed information on this chapter see also the FlexFrame™ for SAP® Version 5.0A manual

"FA Suite V9.0: myAMC.FA_Agents - Installation and Administration".

8.1 FlexFrame Event Management Expert

To view the trap in the virtual, physical and service context FlexFrame provides

- standardized traps with events
- pool and group information
- the virtual and physical server and service information

The events are based on the myAMC.FA_Messenger service and part of every FlexFrame installation since FlexFrame Version 3.1. The standardized FlexFrame trap format and the FlexFrame MIB for the trap format make it easy to integrate in every management platform.

In addition filtering and routing of traps is possible. The myAMC.FA_Messenger provides a standard myAMC.FA Mib structure, this is the base for the fast and easy integration in Enterprise IT Management products, one of the additional values is the frame id to identify FlexFrame Events very easy in an enterprise IT console.

The event management option consists of log file agent to produce SNMP traps from log files e.g. the SAP CCMS log files and to integrate it into the FlexFrame IT management environment with a central IT Data Warehouse (Repository, ITDW).

For detailed information on this chapter see also the FlexFrame™ for SAP® Version 5.0 manuals

“FA Suite V9.0: myAMC.FA_Messenger - Installation and Administration”,

“FA Suite V9.0: myAMC.FA_LogAgent - Installation and Administration”.

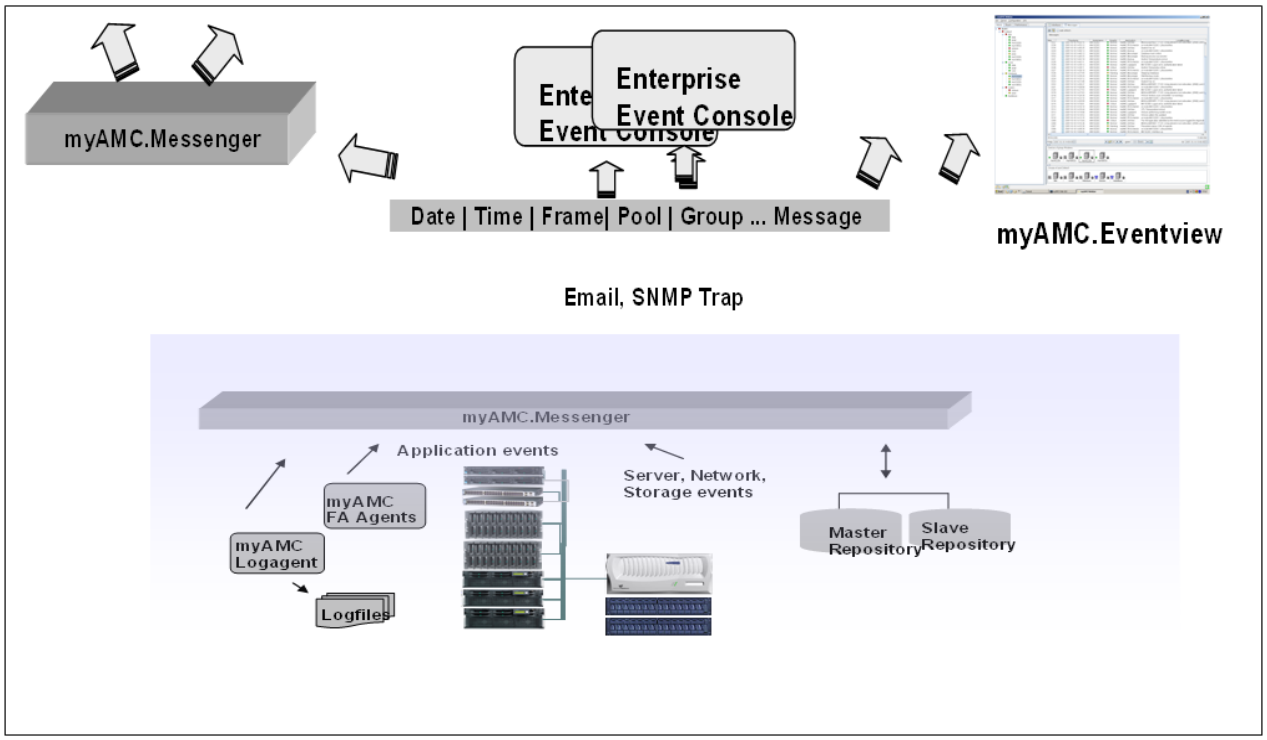

8.2 Performance and Capacity Management

The FA Agents provide optional performance and accounting data. The agents collect node-based, service-based and group-based information. The performance management meets the need for monitoring and historical trending and support intelligent orchestration of resources.

The FlexFrame performance and accounting option requires the activation of additional services on the Control Node. This service does a performance and accounting calculation of the raw data.

The following graphic shows the architecture of the performance and accounting option.

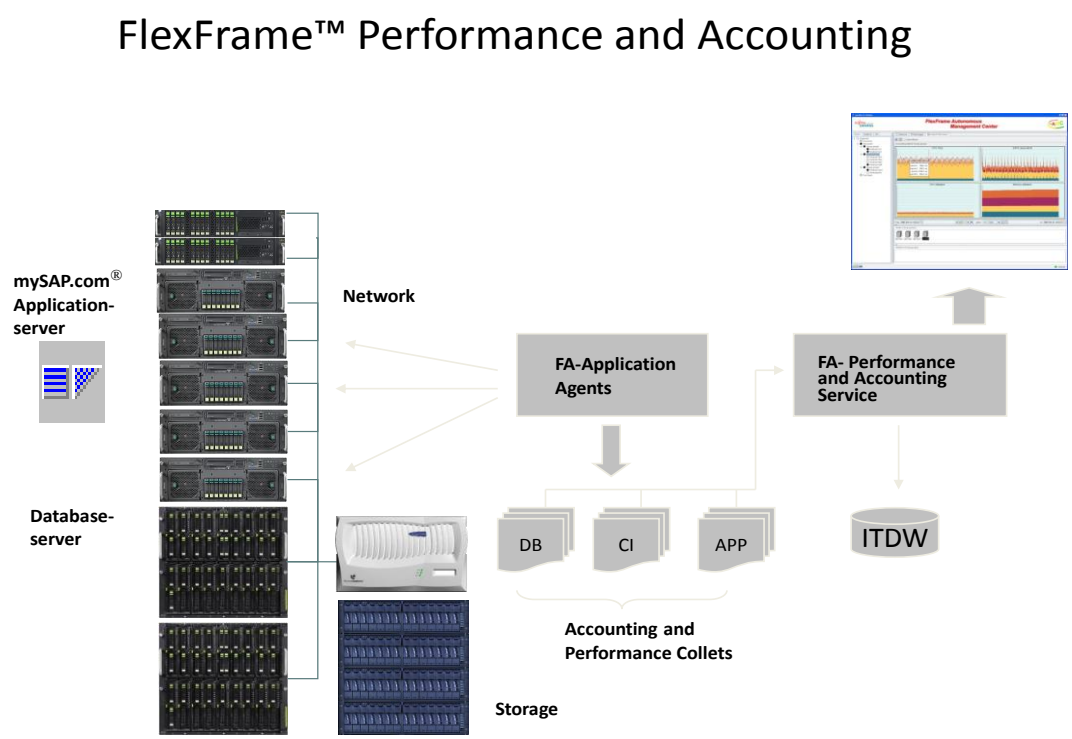

1

8.3 Accounting Management

The accounting option is, like the performance option, an optionally activated part of the FA Agents. The production of the accounting data is a multistage process determining accounting data through aggregation and analysis of the recorded raw data. The accounting option is available for performance values of the server and services and also for the fileutilization accounting.

The performance option enables monitoring and evaluation of the server and services over a longer period of time. For every node the following data are available as a minimal, average and maximum value:

- load of SAP, database or generic services
- other services
- Machine idle

The data of the performance and accounting option can be directly visualized with the FlexFrame FA Web GUI with performance and accounting management plug in. The granularity of the view and the time span can be freely defined.

System SID	Service	Hostname	Time-stamp	CPU ms	CPU %	Mem Kb	Mem %	SAPS abs	SAPS %
P22	DB	Host 1							
P22	CI	Host 2							
P22	APP	Host 3							
P22	J	Host 4							
P22	JC	Host 5							
P22	SCS	Host 6							
P22	ASCS	Host 7							
	Backup	Host 1							
	xy	Host 3							

per Report-cycle **Min, Max, Avg, Total**

Fig. FlexFrame service based accounting option

Important parameters for the accounting are detection and report cycles. The detection cycle defines the number of measurements within a report cycle. The minimum, maximum and average values are calculated on the basis of individual measurements for a report cycle. The detection cycle therefore always corresponds to the detection cycle of the FA Agents, which is also a parameter for the autonomy function.

In addition to the accounting data the FlexFrame IT management option allows you to transform your accounting data directly to billing data. A flexible transformation table allows you compute directly the billing value. The billing factor can change independency from service contract number, system ID, service type and day or time ranges. So you get a flexible out of the box billing model absolutely integrated in the FlexFrame environment with no extra effort.

CPU/ SAPS Values

Aggregation cycle

Service-Contract	System ID	Service typ	fromDate	toDate	dayType	from Time	to Time	Service level rule	Accounting rule	Accounting Price	Unit
SC_12345	P22	all	01.01.1900	01.01.2100	workday	00:00	24:00	Sapsrule	Standard	0.25 €	
SC_12345	P22	all	01.01.1900	01.01.2100	weekend	00:00	24:00	Sapsrule	Standard	0.15€	
SC_12345	P23	DB	01.01.2006	01.01.2100	daily	00:00	24:00	Sapsrule	Standard	0.30€	
SC_12345	P23	other	01.01.2006	01.01.2100	daily	00:00	24:00	Sapsrule	Standard	0.20€	
SC_12345	Q22	all	01.01.2006	01.01.2100	daily	00:00	24:00	Sapsrule	Standard	0.15€	

Accounting cycle

8.3.1 Filesystem Monitoring

A new feature together with the Performance, Capacity and Accounting option of the FA-Agents is the filesystem monitoring. The FA-Agents are able to monitor the filesystem utilization for configured filesystems. For each filesystem individual thresholds can be defined. On transgression of the limits an event is generated and can be dispatched with the myAMC.FA_Messenger expert functionality.

8.3.2 Performance, Capacity and Accounting Profiles

The Performance, Capacity and Accounting collects statistics on CPU and SAPs consumption calculated on server, system or service level to provide powerful analysis and reporting capabilities. An additional feature is to define profiles in which limiting values for the CPU and SAPs consumption can be defined. On transgression of the limiting values a myAMC.FA performance event is generated.

The profiles can be defined on servicelevel, systemlevel and nodelevel. The user can individually decide which profiles exist and for which values events have to be generated. A profile can apply to all systems, nodes and services or pool-, system- , service-, group- or node-specific, depending on the hierarchylevel of the profile.

8.3.3 Dynamic Workload Management

The Performance profile together with the myAMC.FA_Messenger can be used to establish a dynamic workload management. The average CPU load event can activate a reaction on the myAMC.FA_Messenger to start an additional application server.

8.4 FlexFrame Filesystem Utilization and Filespace Accounting Option

This option is available as of version 8.0 of the FA-Agents and enables the observation of filesystems and the node- and system- or service-based accounting of the filespaces. The fill degree of filesystems on local servers and also on storage systems (NAS and SAN) can be supervised through the filespace accounting system. On reaching a specific percentual threshold this system will generate an event. The second feature of the filespace accounting option is to store values for the filespace accounting in the IT Datawarehouse repository.

The following graphic shows the architecture of the filesystem utilization:

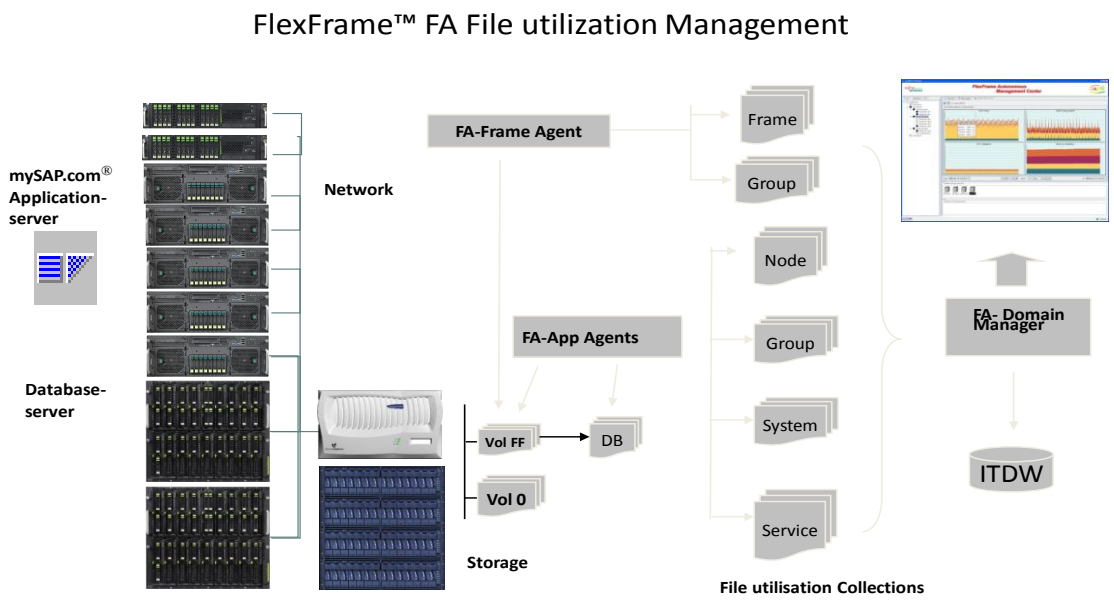

1

The data for filesystem utilization and filespace accounting can be determined with the FA-App-Agents and the FA-Frame-agents. The parameter- and rulesets for the filesystem utilization allow very granular definitions, in which cycles and for which filesystem or file the storage consumption and the increase will be ascertained.

Statistic definition in the ruleset

The observed filesystems, directories filesan be staticly defined in the ruleset i.e. the start of the FA-Agents activates the collection of the data and the calculation for every reportcycle.

Dynamic definition in the ruleset

The ascertainment of the filesystems, directories and filedata can also be achieved in dependency to systems and/or specific servicetypes. The generation of the filesystem utilization data ensures only if the defined system or a specific servicetype actually runs on a node. The FA-App-Agents ascertain through the dynamic triggering.

8.5 Management Reporting Option

Based on a crystal reports technology with integrated viewer, predefined report templates can be used directly to report event, performance and accounting data.

9 Virtualization of SAP Applications

Before the concept of FlexFrame was created, it was common practice to install SAP and database services on a "fixed" server.

Usually one IP address was used for:

- the host itself
- database
- each SAP Instance

Once installed, the service could not easily "be moved away" from its host.

In a cluster configuration, virtual IP addresses were introduced to be able to switch a service from one host to another. FlexFrame is using this technology for virtualization. This concept is one of the fundamental pillars enabling FlexFrame to constitute as a Dynamic Infrastructure Solution.

9.1 SAP Applications Managed by SAP ACC

Since FlexFrame Version 3.0, FlexFrame is compliant with the SAP Adaptive Computing Controller (ACC), see chapter SAP ACI – Certificate of Compliance.

With the full integration of SAP ACC, FlexFrame provides the standardized virtualization layer based on SAP NetWeaver™, this means the virtualization of all released SAP services, the use of SAP's standard user interface to manage SAP services with additional benefits (task planner & mass operation e.g. start or stop a complete landscape with one click), and the integration in SAP's admin & support infrastructure (SAP Solution Manager, NetWeaver Administrator).

The ACC provides a "userexit" which is used to inform the FlexFrame Autonomous Agents (FA Agents). In this way the ACC is integrated into the environment of Control and Application Agents. While the ACC provides a central console for the SAP instances, FlexFrame Autonomous Agents keep services and nodes highly available and execute self-repair functions.

The use of SAP ACC is an optional FlexFrame feature.

10 Pools and Groups

FlexFrame offers advanced functions for partitioning a FlexFrame environment into service-specific or customer-specific server pools and groups. This may be interesting for large installations or application service providers.

10.1 Server Pools

A pool is a number of Application Nodes belonging to the same department or a customer with exclusive hardware requirements. FlexFrame systems can be divided into pools. Each FlexFrame system consists of at least one pool. In a pool, all servers may communicate with each other, but not with the servers of other pools.

Servers of different pools can use different copies of the OS.

Server Pools are separated into different network segments. Server LAN, Storage LAN and Client LAN are pool specific.

The Control LAN is connecting the Control Nodes with all servers of all server pools.

All server pools share the Control Nodes, the Control LAN segment and the NAS storage containing the common volFF volume.

10.2 Server Groups

Within a server pool, various types of hardware can be used with different characteristics, such as operating system, architecture, number of CPUs and RAM size. The bulk of servers can be divided into groups of servers with similar operating systems and hardware performance. Servers in a single group have to be able to take over the services of each other server in this group. This ability may be very useful for groups of high-performance database servers or groups of medium-performance application servers. Each pool consists of at least one group. In FlexFrame groups are denoted by server groups or pool groups.

Each SAP application running in a pool can use one or more servers in one or more groups of servers in the same pool. Each instance of this application runs in a selected group. In case of a failure, switchover to another server is possible in the same group. For example, a set of servers can be divided into high-performance database servers and smaller application servers. The group configuration makes sure that database instances run in a group of high-performance servers, while application instances stay on groups of smaller servers without interfering with each other.

All servers in a pool share the same LAN segments, even if they belong to different groups. They do not share LAN segments with other pools, except the Control LAN. If SAN storage is to be used for SAP databases, the configuration can be built up in a way, that all servers of a group share the same SAN volumes.

11 Prerequisites and Assumptions

Certain prerequisites and assumptions apply for the FlexFrame concept, and these have to be taken into consideration.

11.1 Computing Center Infrastructure

FlexFrame is based on an existing computing center infrastructure. A stable air conditioning is expected. The necessary cooling capacity is depending on HW configuration.

An uninterruptible electric power supply and at least for the Control Nodes two independent current entries are highly recommended.

11.2 Physical Host Names and IP Addresses

Physical host names and their correspondent IP addresses are used by FlexFrame system software for addressing physical nodes over a selected LAN segment. Physical host names may not be used by SAP applications. Host names are differentiated by a suffix to permit an Application Node to be addressed via a selected LAN segment:

- For the Storage LAN this suffix is `-st`
- for the Server LAN `-se`
- for the Control LAN `-co`
- for the Client LAN no suffix is used.

These host names are not restricted by SAP conventions like virtual host names

11.3 Virtual Host Names and IP Addresses

The virtual host name and its corresponding virtual IP address are used by applications for addressing the server on which an SAP instance is running. This virtual IP address will be assigned to a physical host before an SAP instance is started. The virtual IP address thus identifies the actual host on which the SAP instance is actual running, and if an SAP instance is moved to another server this SAP instance is always addressable using the same virtual IP address.

Virtual host IP addresses are also used for communication from applications outside FlexFrame systems, i.e. SAP front ends such as SAPGUI.

Virtual host names and virtual IP addresses are generated automatically by the FlexFrame Management Tool (former Planning Tool), that supports the planning and configuration of a FlexFrame environment.

Formation rules for virtual host names

The virtual host name is formed from the service type, the two-digit instance number for the service type `app` and the SID:

```
<service_type>[<ID>]<SID>[<LAN_type>]
```

`<service_type>` can be one of:

```
db    - database instance
ci    - central instance (ABAP)
app   - application instance (ABAP)
scs   - SAP central (common) services (JAVA)
ascs  - SAP central (common) services (ABAP)
jc    - JAVA central instance
j     - JAVA application instance
lc    - Live cache instance
ers   - Enqueue Replicated Server instance
```

`<ID>` is a number from 00 to 96 (except 2, 25, 43, 72, 89) for `app` and `j` only. It is empty for other service types.

`<SID>` is the system ID of a SAP system.

`<LAN_type>` can be one of:

```
-se             Server LAN
empty string          Client LAN
```

This host name formation rule for virtual services is mandatory for the FlexFrame infrastructure solution. Some components rely on this rule.

In a FlexFrame environment, each node name must be unique. However, each node may have multiple host names that are derived from the node name using a defined naming rule.

In the SAP environment, host names are currently limited to 13 alphanumeric characters including the hyphen (“-“). The first character must be a letter. In the SAP environment host names are case-sensitive (see SAP Note No. 611361).

11.4 Access Rights

The current implementation of the NFS client does not allow specification of the client's IP address; hence the virtual addresses cannot be used for NFS exports (which control access to data paths). Therefore any client in a specific pool can and must be able to mount any path of this pool.

This condition is valid within each server pool but is not permitted between different server pools.

12 FlexFrame for SAP Backup Solution for Oracle with NetWorker

The FlexFrame Backup Solution for Oracle is based on the backup software products

- NetWorker
- NetWorker Module PLUS for Oracle
- NetWorker Module PLUS for Oracle for NDMP NetApp

It provides the following features:

- **Redo log Backup**:
  Redo logs are automatically copied to backup media and removed from disk according to user defined policies

- **Integrated User Interface for Recovery**:
  The recovery user interface integrates database restore and redo log recoveries. Recover options are Crash Recover, Point-in-Time Recover and Offline Recover. In case of NAS storage with NetApp snapshot- and NDMP restores are also integrated.

- **Integration in FlexFrame Virtual Service concept**:
  When a database instance fails over to another Application Node, the respective NetWorker components are automatically switched as well.

Additional features in the case of NAS storage with NetApp:

- **Snapshot Backup**:
  Snapshot backups take only seconds and don’t have any performance impact on storage and Application Nodes. This allows frequent backups in short intervals even during working hours. Up to 255 snapshots per volume can be stored in parallel. Recover from a snapshot takes only minutes, depending on the recover scenario. Time required to restore redo logs can be significantly reduced by snapshots taken in short intervals.

- **NDMP Backup**:
  NDMP (Network Data Management Protocol) is the standard protocol for the backup of NAS systems. NDMP backup provides server-free backup since the NDMP backup goes from the NAS device via Fibre Channel (FC) or SCSI to the tape drive without any impact on the Application Node. NDMP backup to tape protects against severe hardware failures and against a disaster.

  During backup with NetWorker Module PLUS for Oracle for NDMP NetApp, a snapshot of the entire volume is made, but only files of DBs (SIDs) that are set in backup mode, are written to tape. Data can be restored from tape or snapshot SID-specifically. Also, only single data files can be restored.

A corresponding Best Practice Paper can be found at

https://partners.ts.fujitsu.com/com/products/infrastruc-solutions/FlexFrame .

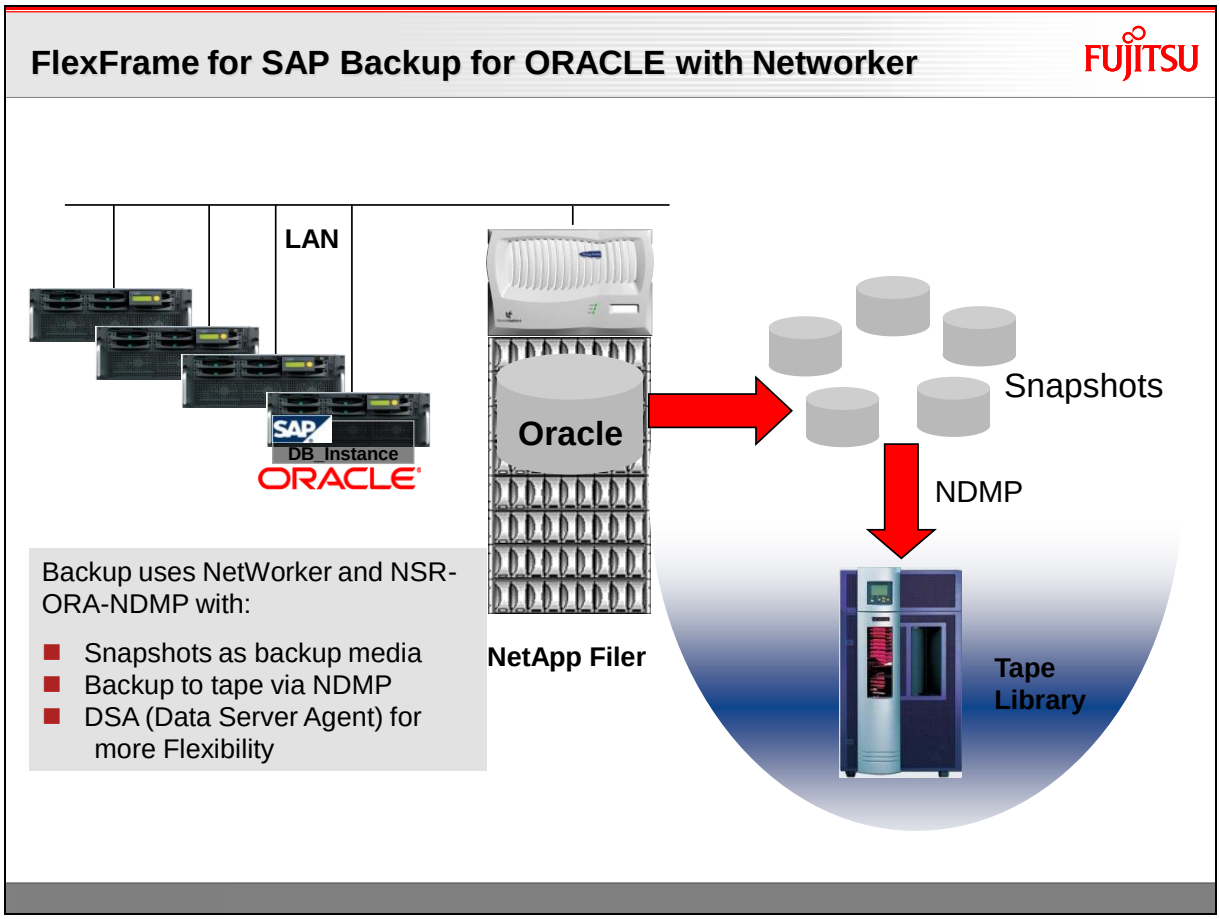

12.1 Backup Hardware

The FlexFrame for SAP backup solution for Oracle adds a dedicated backup server, a tape library and a SAN infrastructure to the FlexFrame environment. An important benefit of dedicated NetWorker server hardware is the ability of a fast recovery from severe hardware failures of the FlexFrame system. For backup of clustered NAS a FC connection to tape drives is necessary.

The following shows a list of proposed hardware components:

- A PRIMERGY server with Linux as backup server. For reasons of high availability two of these servers may be configured as a backup server cluster.
- A tape library (e.g. Scalar i500 with LTO 3) drives for automation of the tape handling.
- A FC environment consisting of a 16 port Silkworm switch with multi mode FC GBICs and cabling for interconnection of the FlexFrame and the backup environment.

12.2 Backup Software

FlexFrame Backup for Oracle environment consists of the following software components:

- A NetWorker server software release 7.4 for Linux. This component is responsible for administration of all backups by controlling the robot and communicating with the NAS system via NDMP.
- A NetWorker Client release V7.4 is running on each Application Node.

- "The NetWorker Module PLUS for Oracle" release 5.1 is running on each virtual Oracle database server along with the Archive-Monitor for redologs release 2.1.

12.3 Other Backup Solutions

The deployment of other backup solutions in a FlexFrame™ for SAP® environment assumes that the solution was tested and certified with FlexFrame for SAP (see chapter Third-Party Software on page 30).

13 Third-Party Software

The interfaces FlexFrame provides are basically the standard interfaces of the operating systems supported with FlexFrame. The Operating Systems are installed in a way that they behave like standard OS installations for SAP. One big difference is that the root file system "/" is mounted and accessed via NFS. The third-party software has to be capable to deal with NFS-only file systems.

Some third-party software (e.g. system management tools, backup tools, etc.) may not be able to handle moving (virtual) services and assume a fixed server-to-service relationship. In addition, those products may write files at locations that are mounted read-only and must therefore be relocated to a read/write location.

During implementation of a FlexFrame project the required third-party products must be identified and a project-specific analysis of this product must be initiated. Fujitsu offers such validation of third-party software as a service.

As mentioned in 6.5 it is also possible with the FlexFrame Autonomous agents from FlexFrame for SAP to integrate "third-party services" into the autonomy scenarios. For this purpose a generic service is defined through a set of parameters which are used for its identification and which generate the service states.

14 Support

14.1 Fujitsu Customer Support

FlexFrame environments are fully integrated into the support structure of Fujitsu. A large number of consultants are available for design, planning, installation and further customer support.

14.2 SAP

The FlexFrame infrastructure solution approach is well-known and fully supported by SAP. Customers can create problem reports through existing SAP mechanisms (SAP Support Portal) in the same way as for conventional SAP system environments. Fujitsu are fully integrated and are informed by SAP if a problem is platform-specific. Support staff from Fujitsu work on such problems using SAP's support structures.

During installation of a SAP environment (FlexFrame or not) it is advisable to open TELNET access and a SAP Support Portal connection to the installed system. If a problem occurs, the customer can open the SAP Support Portal connection (with TELNET or SAPGUI) and provide logon information for support purposes. This remote support allows rapid detection and correction of problems.

15 Benefits of the FlexFrame Platform

The following sections briefly describe some of the benefits of the FlexFrame concept.

15.1 Shared Use of Storage Space

Legacy installations require dedicated LUNs for data volumes. Since growth of the data volumes is hard to predict, a lot of free space is required to assure continuous operation of the database.

With the NAS systems as storage, multiple databases (SIDs) can share the same space (volume), and the disk space is utilized much more efficiently.

15.2 Shared Usage of Servers

Due to the virtualization of SAP services, the total number of servers required can be reduced. Following are two sample scenarios.

15.2.1 Less Servers Required for High Availability

Legacy concepts usually require a pair of dedicated servers for clustering SAP databases or central instances. In a larger SAP environment, multiple cluster pairs of this type lead to a large number of servers just to cover hardware failures. Usually the capacity installed is much larger than the capacity used.

In FlexFrame, spare servers are used to take over failed services from the complete FlexFrame environment, or, if server pools and groups are used, from the complete pool/group, or in case of pool independent spare servers even from the complete FlexFrame server park.

15.2.2 Better Coverage of Performance Peaks

In larger SAP environments, multiple SAP systems (SIDs) communicate with each other. This may result in the situation where an ERP system processes, for example, month-end calculations and passes on information to a BW system. This BW system now builds aggregates of the new information. The servers for both the ERP and BW systems must be sized to be able to handle the maximum load. In legacy system environments both servers may be idle most of the time.

With FlexFrame, the available hardware can be shared between processes. Processes can be moved to a faster host when the load increases, and moved back to slower hardware after a load peak has been processed. The number of high-end servers that are able to handle maximum peaks can therefore be significantly reduced.

15.3 Less Administration

The shared operating system feature provides advantages in an administrator's daily life.

Linux Application Nodes share their "/" root file system. Hence, updates of software packages are available immediately to all Application Nodes. This means that, after updating and testing with a Test Application Node, you can activate a new root file system version by simply rebooting the Application Nodes.

15.4 Less Downtime During a Change of OS

Imagine you want to update a large number of servers to a new version of Linux. Usually this means creating backups of each server, and updating and testing each server individually.

With FlexFrame you can simply create a new directory on the NAS storage and install your new version of the OS in this directory.

Once tested sufficiently, you only change the boot configuration to use the new root file system, and you reboot the Application Nodes whenever convenient. And there is a most simple way of fallback to the old OS version if needed.

15.5 Improved Change Management

An improved change management in FlexFrame for SAP environments is guaranteed by

- dynamic SAP operation - you have a very high level of adaptivity and short response times in fulfilling requirements
- the higher quality of service - possible errors in making changes are minimized and downtimes are reduced
- low-cost adaptations - personnel expenses and costs for external service providers and capital spending are reduced

- Simple further development of the SAP environment - the standardization of the IT infrastructure minimizes the time and work involved in adaptations, FlexFrame for SAP scales very easily and virtually without limits and systems and new services can be rolled out at little cost and effort.

15.6 Simplified Multi-Client Operation

If services are provided for multiple customers, a conventional approach requires a separate infrastructure for each one. As a result, the topics and problems in planning, implementation, operation and adaptation multiply. Now FlexFrame is the ideal platform for supporting multiple clients with a single IT infrastructure or running different SAP systems. The various environments are kept strictly apart from each other by separate networks. It is also possible to reassign hardware between different customers or systems in a matter of minutes. Numerous functionalities make it possible to configure SLAs for the individual customers to handle individual agreements. Functions such as event and error management, accounting management, performance and capacity management and convenient reporting options additionally support operations management for various customers.

15.7 No Oracle Standby Database Server Required

A shadow database is commonly used for an additional level of high availability. An additional server is required for such a concept. The changes to the database ("redo logs") are shipped to this remote server and are applied to a copy of the database with a defined time lag. That means the shadow database is always some hours behind its master DB. If a logical error occurs in the original database, the shipping of the logs is stopped and the "waiting" logs will not be applied. If the error occurred before the fixed time delay (e.g. three hours), a restore from other media is required. Most decision-making processes take a long time to decide whether to switch over to the standby database or fix the issue in a different way.

Since FlexFrame is based on NAS storage systems that enable large amount of snapshots per volume to be created, a more granular recovery procedure has been provided. It is possible to separate logs and data files by volume. You can make rotating snapshots, e.g. on an hourly basis. If an error occurs the database will be stopped. A snapshot from a time before the error happened is restored only for the data volume (or even on data file level). Since the log files are still available, it is now possible to recover to a time just before the error occurred.

In FlexFrame, the additional server for an Oracle Standby Database is not necessary.

16 FlexFrame and SAP® Business Suite

FlexFrame is a platform explicitly designed for the SAP® Business Suite. The components supported by the FlexFrame infrastructure solution are limited by SAP's product availability matrix (PAM) and the managed systems supported by SAP's ACC.

Components which are not supported by FlexFrame should run on separate, dedicated servers. FlexFrame Autonomous Agents do not monitor such components unless these are virtualized services integrated into the autonomy monitoring and reaction scenarios by means of the **generic services** (see 6.2).

As of FlexFrame for SAP version 4.1A in addition to the hitherto supported databases MaxDB (SAP DB) and Oracle also the IBM database DB2 is supported on SuSE Linux platforms.

17 SAP ACI – Certificate of Compliance

With FlexFrame, Fujitsu are the first to provide an **SAP adaptive computing compliant** technology platform. PRIMERGY servers with NetApp and EMC are supported.

17.1 SAP ACI for PRIMERGY/SLES10/11 and NetApp Storage

CERTIFICATE OF CONFORMITY

SAP AG hereby confirms that

Fujitsu

has been approved to receive this certificate by providing an

Adaptive Computing Compliant

Concept according to the Adaptive Computing Compliance Test.

Certificate No. 2010007
The test has been performed using the following platform and building block combination:

Computing: **PRIMERGY™ Server family with VMware® ESX and OS SuSE® SLES 10/11 based on FlexFrame™ for SAP® 5.0A**
Storage: **NetApp® Fabric-attached storage (FAS) systems**
Control: **SAP NetWeaver Adaptive Computing Controller 7.2**

Walldorf, August 13, 2010

i.A. Stefan Reichart
VP TD LM On Demand & TCO
SAP AG

THE BEST-RUN BUSINESSES RUN SAP

17.2 SAP ACI for PRIMERGY/SLES10/11 and EMC Storage

CERTIFICATE OF CONFORMITY

SAP AG hereby confirms that

Fujitsu

has been approved to receive this certificate by providing an

Adaptive Computing Compliant

Concept according to the Adaptive Computing Compliance Test.

Certificate No. 2010006
The test has been performed using the following platform and building block combination:

Computing: **PRIMERGY™ Server family with VMware® ESX and OS SuSE® SLES 10/11 based on FlexFrame™ for SAP® 5.0A**
Storage: **NAS Storage: EMC® Celerra™ series**
SAN Storage: EMC® Symmetrix® DMX series, CLARiiON® CX series and Fujitsu FibreCAT CX series
Control: **SAP NetWeaver Adaptive Computing Controller 7.2**

Walldorf, August 13, 2010

i.A. Stefan Reichart
VP TD LM On Demand & TCO
SAP AG

THE BEST-RUN BUSINESSES RUN SAP™

CONTACT
FUJITSU [Fujitsu Technology Solutions GmbH]
FTS TSP ES&S SWE SOL

E-mail: michael.wester@ts.fujitsu.com

Website: ts.fujitsu.com

2010-11-10