

LISA

Large Installation System Administration

AAU
15 Marts 2007


Karsten Thygesen
CTO, Partner, Netic A/S
karthy@netic.dk

- Many servers
 - Multiple architectures
 - Complex setup
- Many users/organizations
 - With different goals and privileges
- High requirements to
 - Security
 - Availability
 - Privacy
 - Manageability




- Low expenses
 - Few system administrators
 - Under skilled administrators
 - Low budget for administrative projects (make projects add to infrastructure)
- Low time-to-market
 - High workload
 - Pressure to cut corners (which means more work later on)
 - No focus on monitoring, maintenance, education
 - Preference for turn-key solutions

- They are the center of the Universe
 - Nothing is more important than their current problem/needs
 - IT can't be that difficult – admin is just clueless - underestimate complexity
- 100% availability, 200% performance, no quotas
 - Even announced downtime is not acceptable
 - Central computers are always slow
 - mine at home is faster!



TO HAVE FUN



And do self-education by investigating different tools, languages etc.

- Unquestionable technical skills in *all* aspects of computing
 - Young with 25 years of experience
 - Good communication skills (talk like an user)
 - Good political skills (talk like management)
 - Manager (both economy, strategy and staff)
 - Projectleader, teamplayer, etc
-
- All done with too few staff, to low budget and too high workload – with no signs of stress!



- Meta administration
- Micro Management -> Macro management
- Individual computers -> groups and clusters!
- Server virtualization
- Stateless clients
- Central
 - File/data store
 - access control
 - user id management
 - Backup
 - Monitoring
 - software deployment, configuration management

- Which storage architecture? JBOD, RAID, SAN, NAS, iSCSI
- Many protocols to choose from
 - NFS (version 2, 3 or 4? Windows support?)
 - CIFS (Unix support?)
 - AFS (complex but nice)
 - iSCSI - new
 - Many research projects still going on



Intransa IP 7500

	CIFS	NFS	AFS
Platform	Windows	Unix	Both
Internet ready	Partly	Partly v.4	Yes
Homogen view	No	No	Yes
Works on low bandwith	No	No	Yes
Scales	100's	100's	1000's
Security	AD	V3: trust, v4: gssapi	Kerberos
Non-trust clients	Yes	V3: no, V4: yes	Yes
Reliability	No	No	Replication
Caching	No	Inode	Yes
Complexity	Low	Low	High

- Central File/Data storage
- With many users
 - Impossible to keep track of all groups
 - Users must be able to define their own groups
 - NFS: v3: no, v4: partly, CIFS: partly, AFS: yes
- Central managed firewalls
- Central managed IDS (intrusion detection systems)
- Central identity management



- Which platform to choose?
 - Tivoli, Legato, HSM, Serverless SAN/NAS
 - Amanda, Bacula?
 - Platform support?
- What do we need to backup?
- How long retentiontime?
- NetApp Nearstorage as cache?
- Offsite tape storage and procedures?
- How do users get their lost files back?
- Backup verification?
- Backup of network configurations etc?



- *Really* aim at getting only one monitoring platform?
- Doing monitoring right can be as expensive as the monitored system it self!
- Monitoring is not red/green but statistics, trends and aggregation!
- What about staffing and escalations?
- Disaster plans?



- Automatic tools to:
 - configure systems and maintain that configuration
 - Deploy software and complex configuration files
 - Do simple, repeatable work
 - Self repairing Monitoring
 - Document (intended) state of system



- Central configuration/rule files
- Describes intended state – not how to change
- Adapts to current system
- Built-in secure configuration and file distribution, reporting, tripwire and more
- Works well across different (Unix)platforms
- Limited windows support
- Can be extended to almost anything
- Powerful but also potential dangerous.

```
editfiles:
  solaris::
    { /etc/inet/ntp.conf
      AutoCreate
      AppendIfNoSuchLine "server ntp01.netic.dk"
      AppendIfNoSuchLine "server ntp02.netic.dk"
      AppendIfNoSuchLine "server ntp03.netic.dk"
      DefineClasses "startntp"
    }
  freebsd|linux::
    { /etc/ntp.conf
      AutoCreate
      AppendIfNoSuchLine "server ntp01.netic.dk"
      AppendIfNoSuchLine "server ntp02.netic.dk"
      AppendIfNoSuchLine "server ntp03.netic.dk"
      DefineClasses "startntp"
    }

shellcommands:
  startntp.solaris::
    "/etc/init.d/xntpd start"
  startntp.freebsd::
    "/usr/sbin/ntpd -p /var/run/ntpd.pid"
  startntp.redhat::
    "/etc/init.d/ntpd restart"
  startntp.debian::
    "/etc/init.d/ntp-simple restart"
```


- Copy files/directories ensure permissions, owner, checksum
- Check/create links, files, directories
- Run shell commands (and use output)
- Check/start/stop/signal processes
- Clean /tmp and similar directories
- Many commands to edit text files
- Check resolver, domainname, timezone etc
- Classes defined on OS, kernel, time, date, hostname, domain, user commands, other rules

- No! – but sure helps
- Cfengine also have limitations
- Bootstrapping is hard
- Unix centric, but with some Windows support
- Other projects aims at the same goal but with different methods
 - Radmin (Mac OS-X focused)
 - LCFG (Linux/Solaris focused)
 - Many windows (only) solutions (SMS)



- No single point of failure
- Loadbalancing
- Clustering
- Redundant Net, firewall, switches, routers
- Multiple ISP's (BGP, AS-numbers etc)
- Replicated storage
- UPS and generators
- Redundant cooling
- Multiple server rooms with redundant setup




donner.zeus.com (admin) Logout

Cluster: OK
1450 bytes/sec
Home
Services
Catalogs
Diagnose
Activity
System
I want to:
Help

Traffic Managers


donner
212.44.21.80


gyros
212.44.21.82

Services

Sort by...

 FTP site FTP (21)	 running	  FTP Default
 MX SMTP SMTP (25)	 running	 <div>  Basic SMTP (Private) No relay for external SMTP </div> <div>  Basic SMTP (DMZ) No relay for external SMTP </div> <div>  Basic SMTP (Private) Default </div>
 DNS slaves DNS (53)	 running	  DNS slaves pool Default
 Main web sites HTTP (80)	 running	  HTTP Default
 Secure web site SSL (HTTPS) (443)	 running	  HTTPS Default

Event Log


24/Feb/2005:00:02:57 +0000 INFO:Deleting statd log file 20050125.log


gyros


24/Feb/2005:00:01:16 +0000 INFO:Deleting statd log file 20050125.log


donner

- VMWare ESX and Virtual Center
- Hardware is just a pool of resources
- Built-in intelligence on hardware utilization
- Finally a common (virtual) hardware platform?
- Create new servers by cloning
- Virtual network components
- VMWare Marketplace for standard components

anoka.vmware.com: VMware Management Interface - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Home Search Favorites Media Print View Source

Address <https://anoka.vmware.com/vmware/en/> Go Links »

VMware ESX Server 2.5.0 build-11343 | root@anoka.vmware.com

Status Monitor Memory Options Refresh | Manage Files... | Help | Log Out

Last updated Tue Mar 1 16:44:08 PST 2005

System Summary

5 Minute Average

Physical Processors (2)		Memory (2.0 G)	
Virtual Machines	41 % <div><div></div></div>	Virtual Machines	1.2 G <div><div></div></div>
System Services	10 % <div><div></div></div>	System Services	589.6 M <div><div></div></div>
System Total	51 % <div><div></div></div>	System Total	1.8 G <div><div></div></div>

Virtual Machines (8)

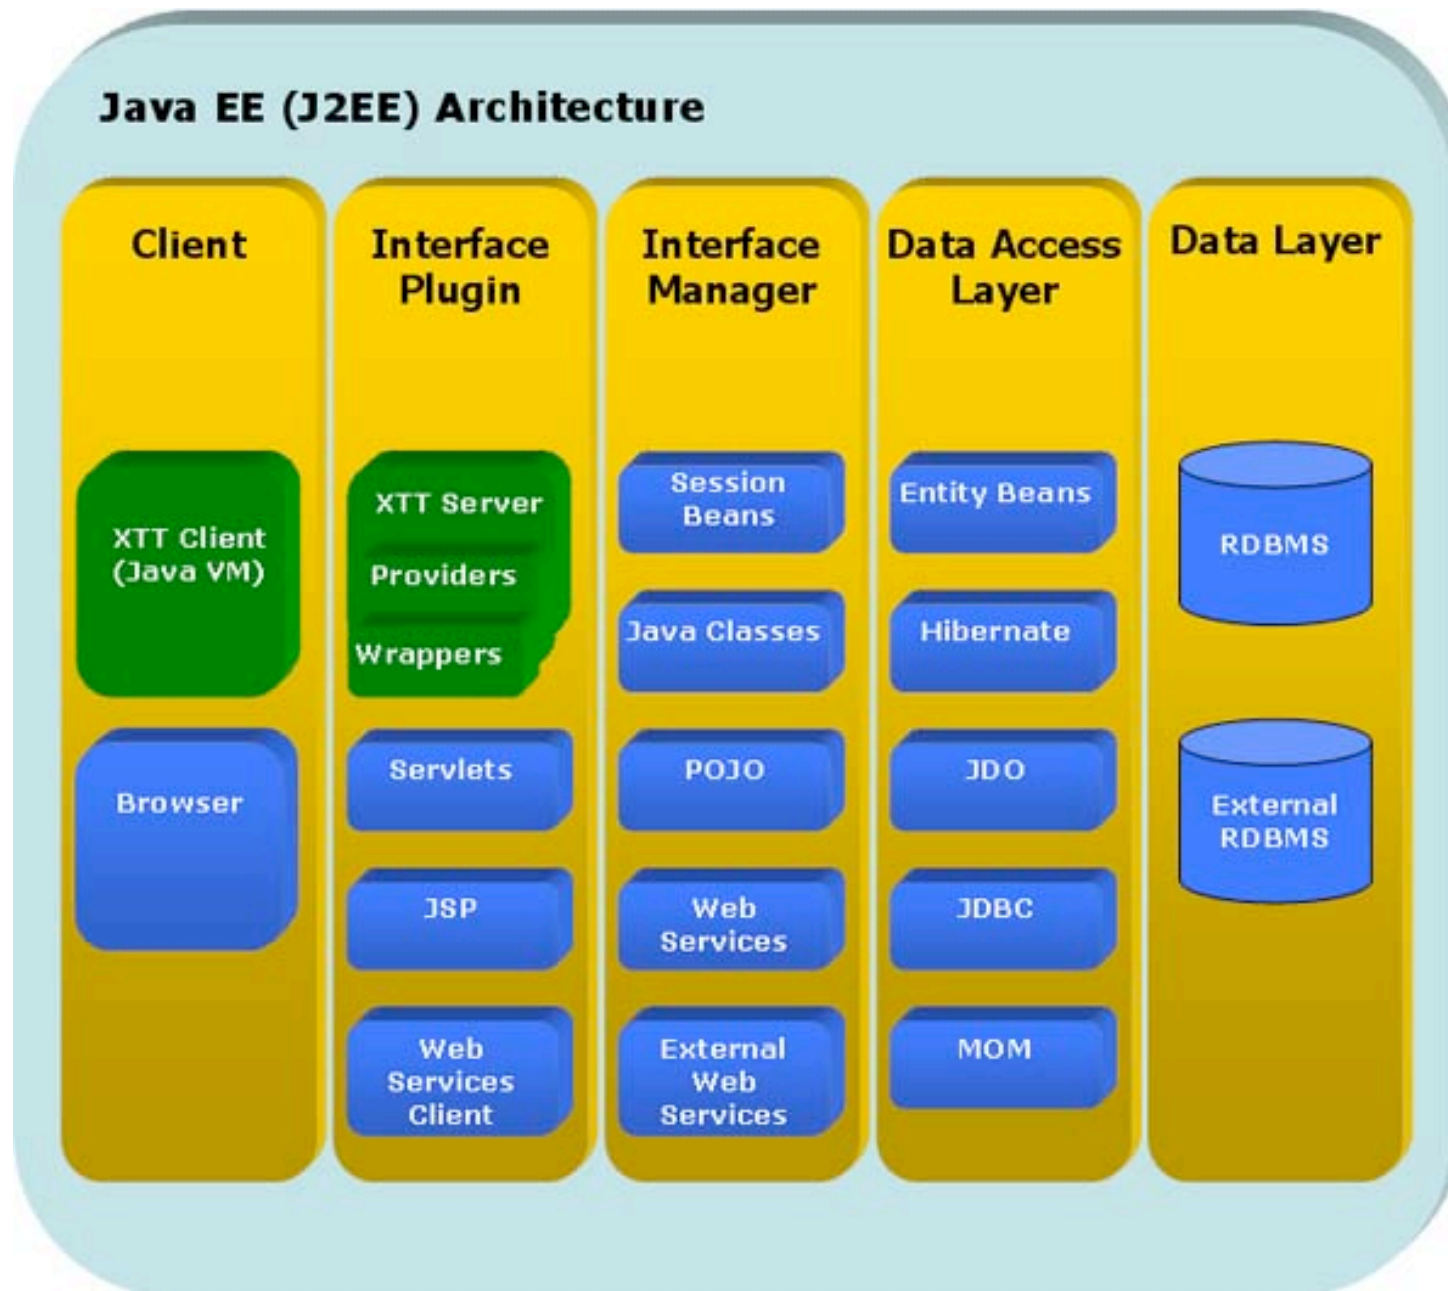
HB	Display Name	Up	No.	% CPU	RAM
	OpenView Operations Powered on PID 11519 VMID 147	5 days	1P	9 <div><div></div></div>	390.8 M
	SiteScope Powered off		1P	<div><div></div></div>	
	W2KDC Powered on PID 11582 VMID 148	5 days	1P	13 <div><div></div></div>	208.6 M
	RHEL Powered off		1P	<div><div></div></div>	
	Lakeside Powered on PID 1048 VMID 154	16 minutes	1P	4 <div><div></div></div>	275.1 M
	SIM42 Powered on PID 11344 VMID 146	5 days	1P	7 <div><div></div></div>	390.5 M
	Novell NetWare 6 Powered off		1P	<div><div></div></div>	
	Director420 Powered on PID 25766 VMID 153	51 minutes	1P	8 <div><div></div></div>	323.8 M

Download VMware Remote Console: [Windows \(exe\)](#) | [Linux \(rpm\)](#) | [Linux \(tar.gz\)](#)

Add Virtual Machine

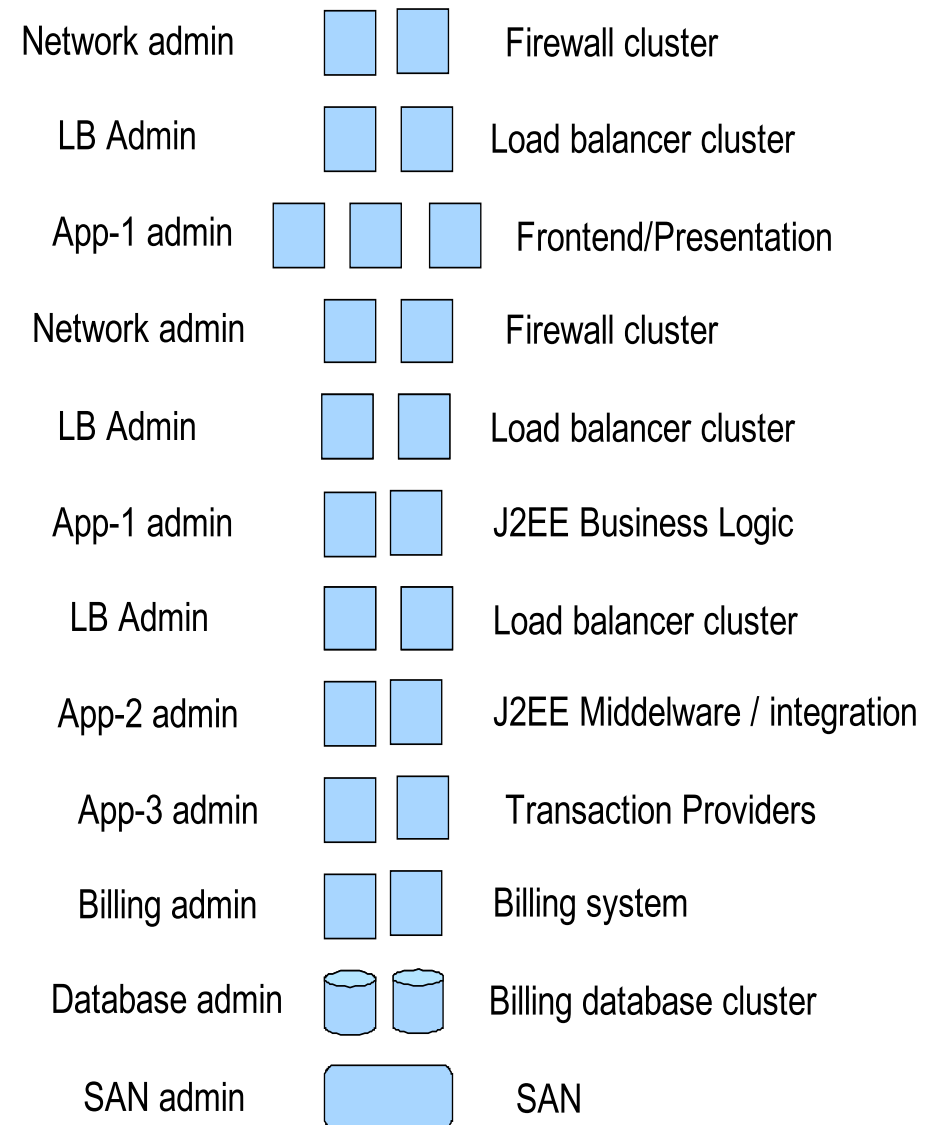
Copyright © 1998-2004 VMware, Inc. All rights reserved.
Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289 and 6,735,601; patents pending.
VMware, the VMware "boxes" logo, GSX Server, ESX Server, Virtual SMP and VMotion are trademarks of VMware, Inc.

Active Internet



- Hardware just a farm of resources
- Heavily virtualization
- SOA – Service Oriented Architecture
- Simpler to write frontend apps, but more complex backend layers
- System administrators needs to have some programming skills and knowledge of SOAP
- Diversion between “system administration” and “application administration”
- Increased need for infrastructure architects

- Many roles
- Many people
- Many layers
- Many servers
- Many data
- Many level of knowledge
- Debugging difficult





- <http://vmware.com>
- <http://usenix.org> (And LISA)
- <http://cfengine.org>

- Study cfengine documentation – perhaps deploy on your own computer
- Study VMWare Server - free download!