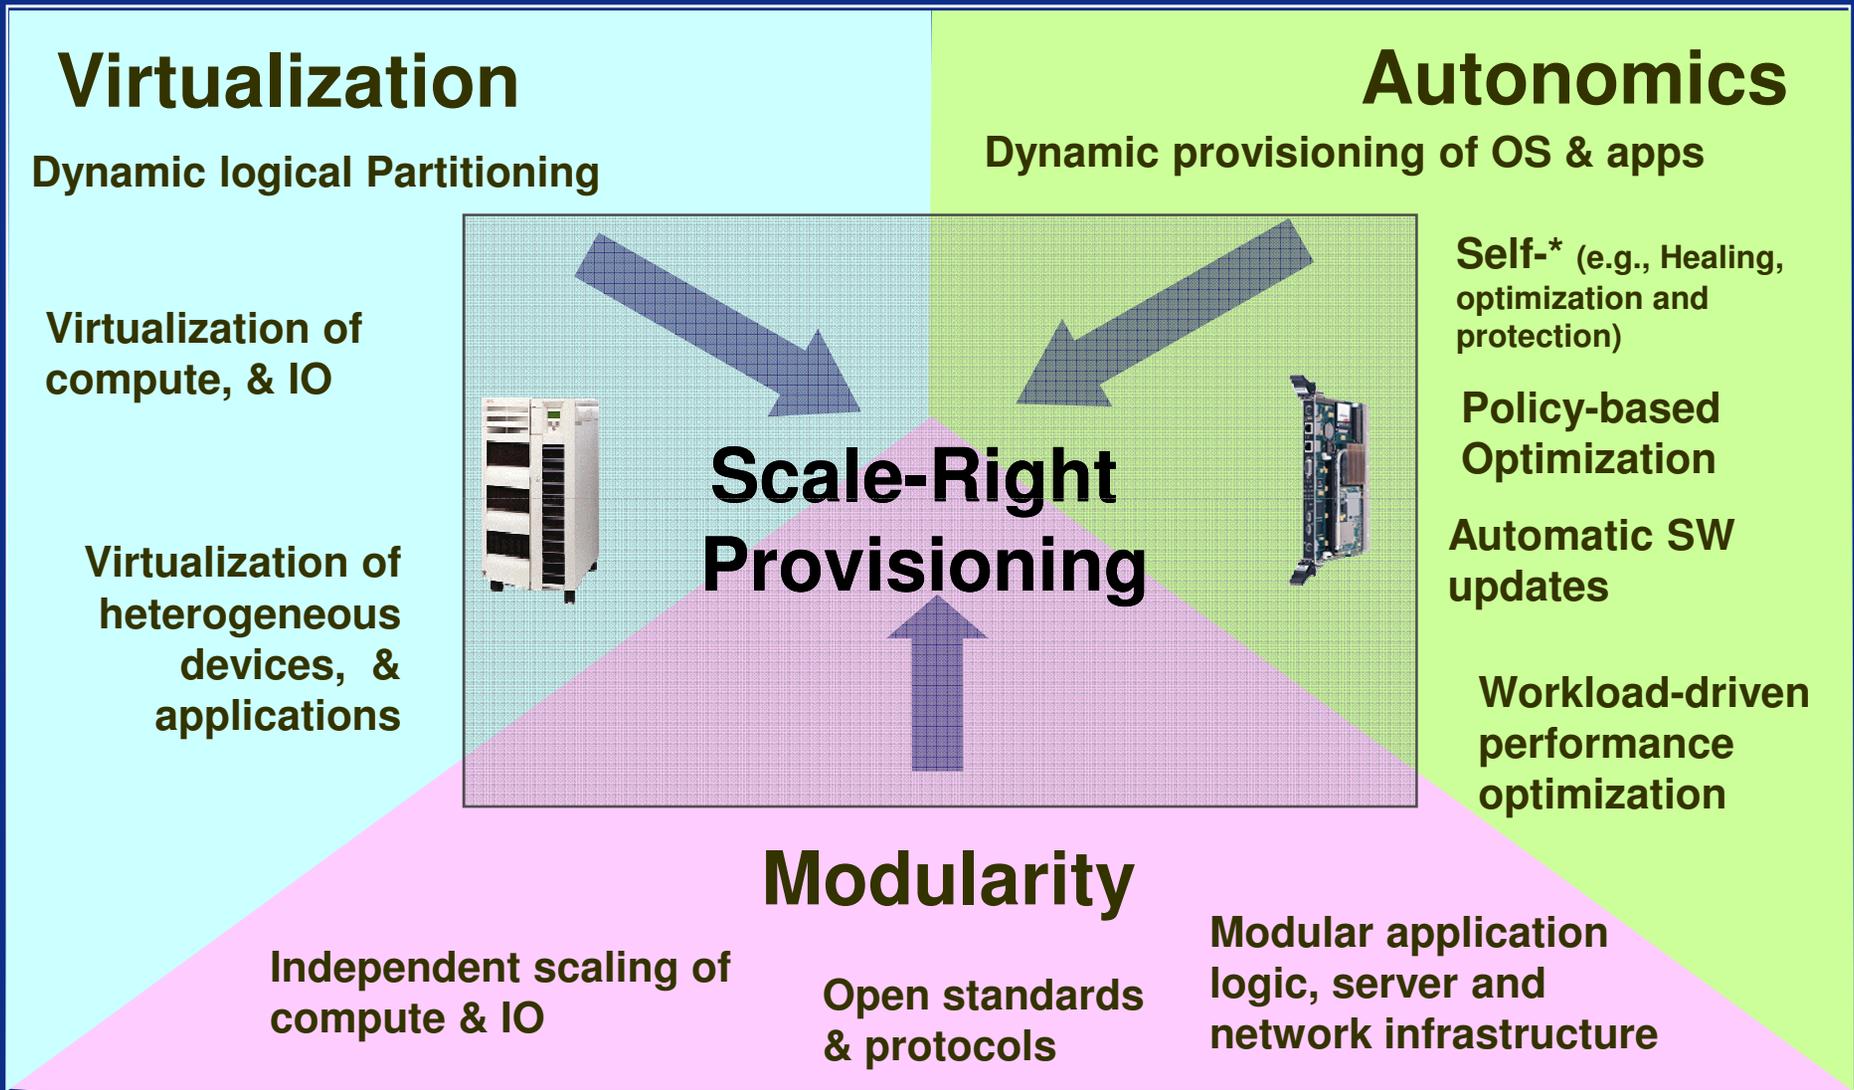


# **Autonomic Computing in Next Generation Datacenters – One Perspective**

**Mazin Yousif, PhD**  
CTO, Avirtec Corporation  
[mazin@avirtec.net](mailto:mazin@avirtec.net)

# Next Generation Datacenters



# Industry and society need for Autonomic Computing

8/12/07: 20K people + 60 planes held at LAX after computer failure prevented customs from screening arrivals

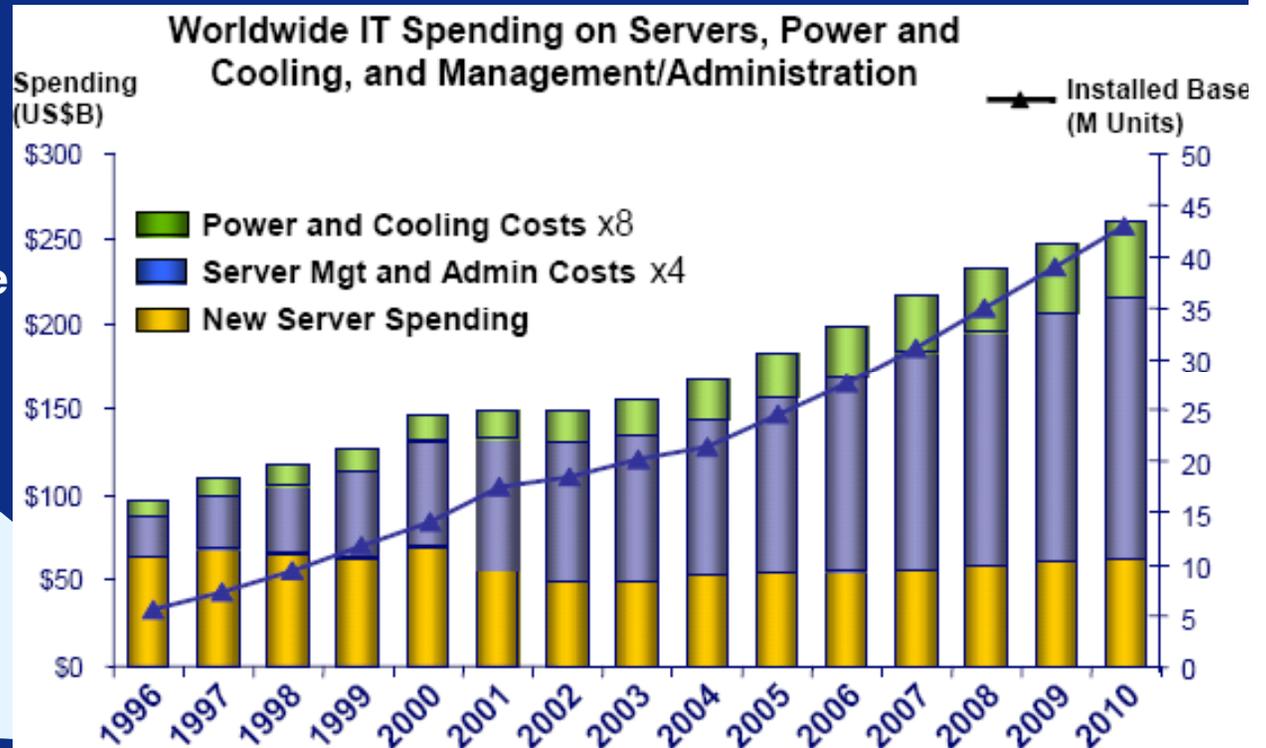
2/27/07: Dow fell 546. Since worst plunge took place after 2:30 pm, trading limits were not activated

8/1/06: UK NHS hit w/ massive computer outage. 72 primary care + 8 acute hospital trusts affected.

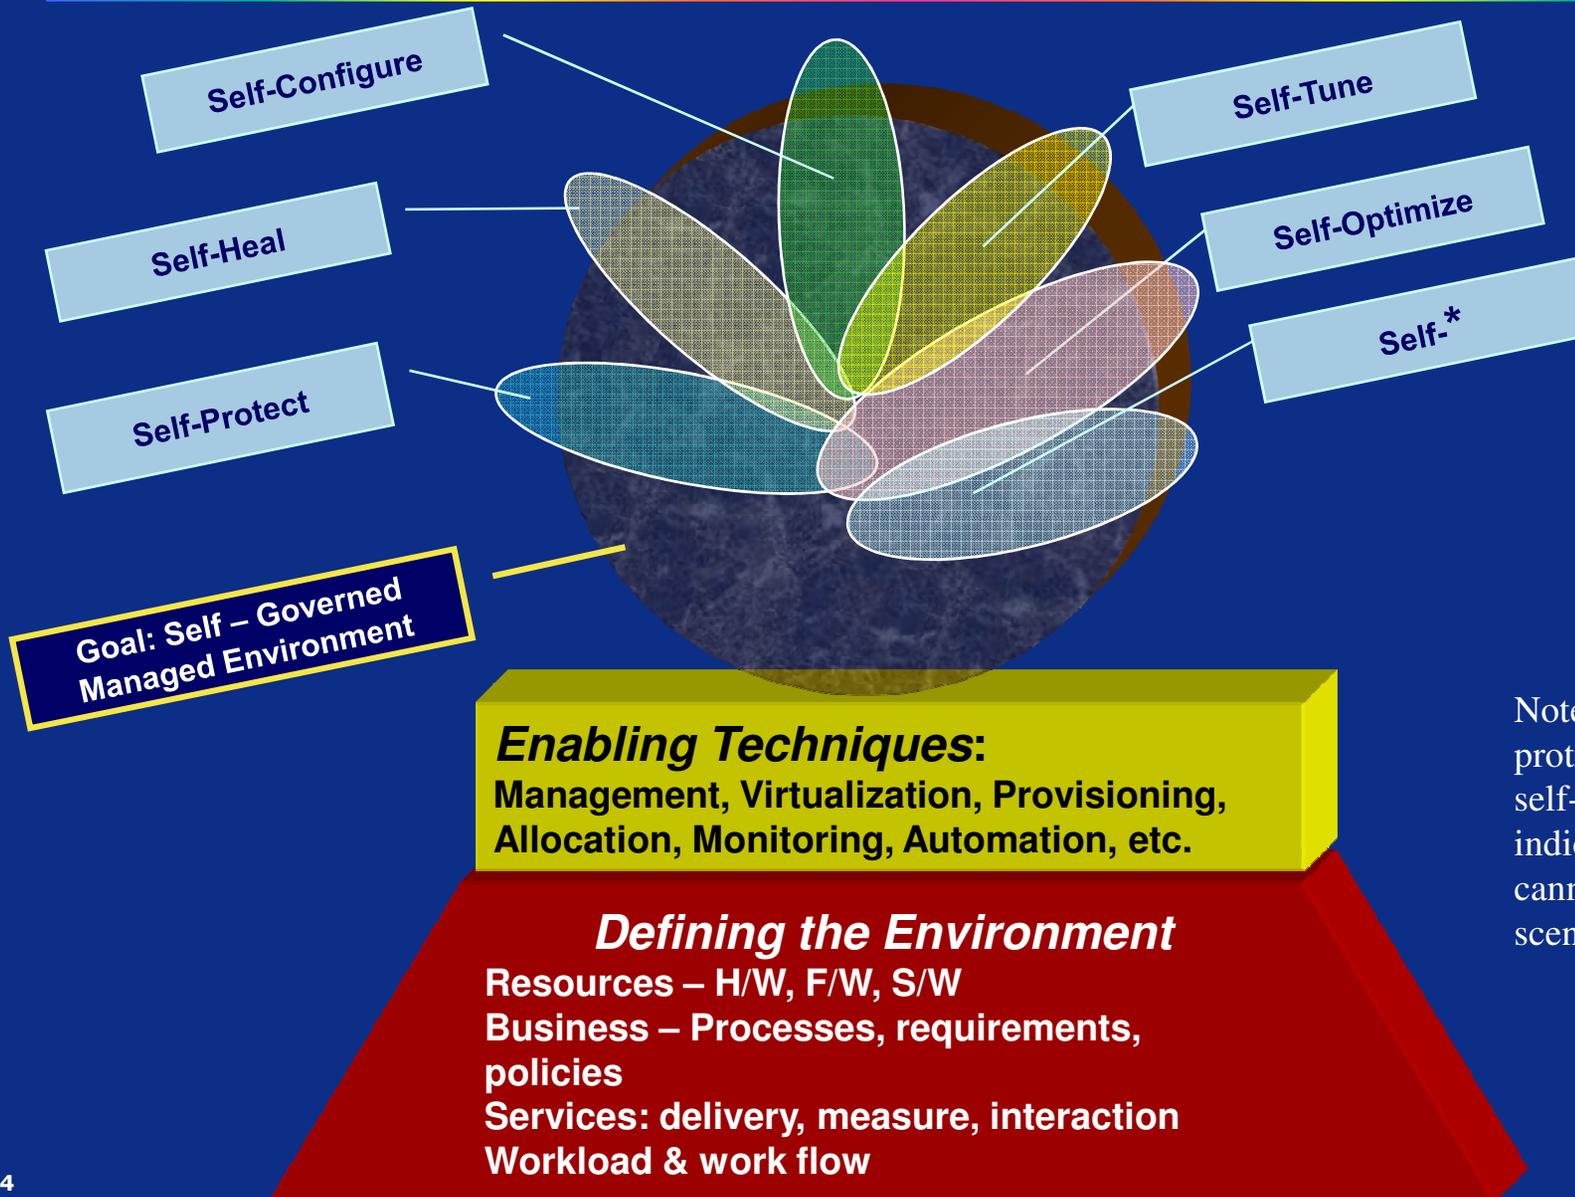
05: 40% firms w/ long computer outage never resume operations. 1/3 out of business in 2 years (Gartner)

1/29/00: NSA "serious computer problem" affected ability to process intelligence for 3 days

8/3/07: (EPA) datacenter energy use by 2011 will cost \$7.4 B, 15 power plants, 15 Gwatts/hour peak



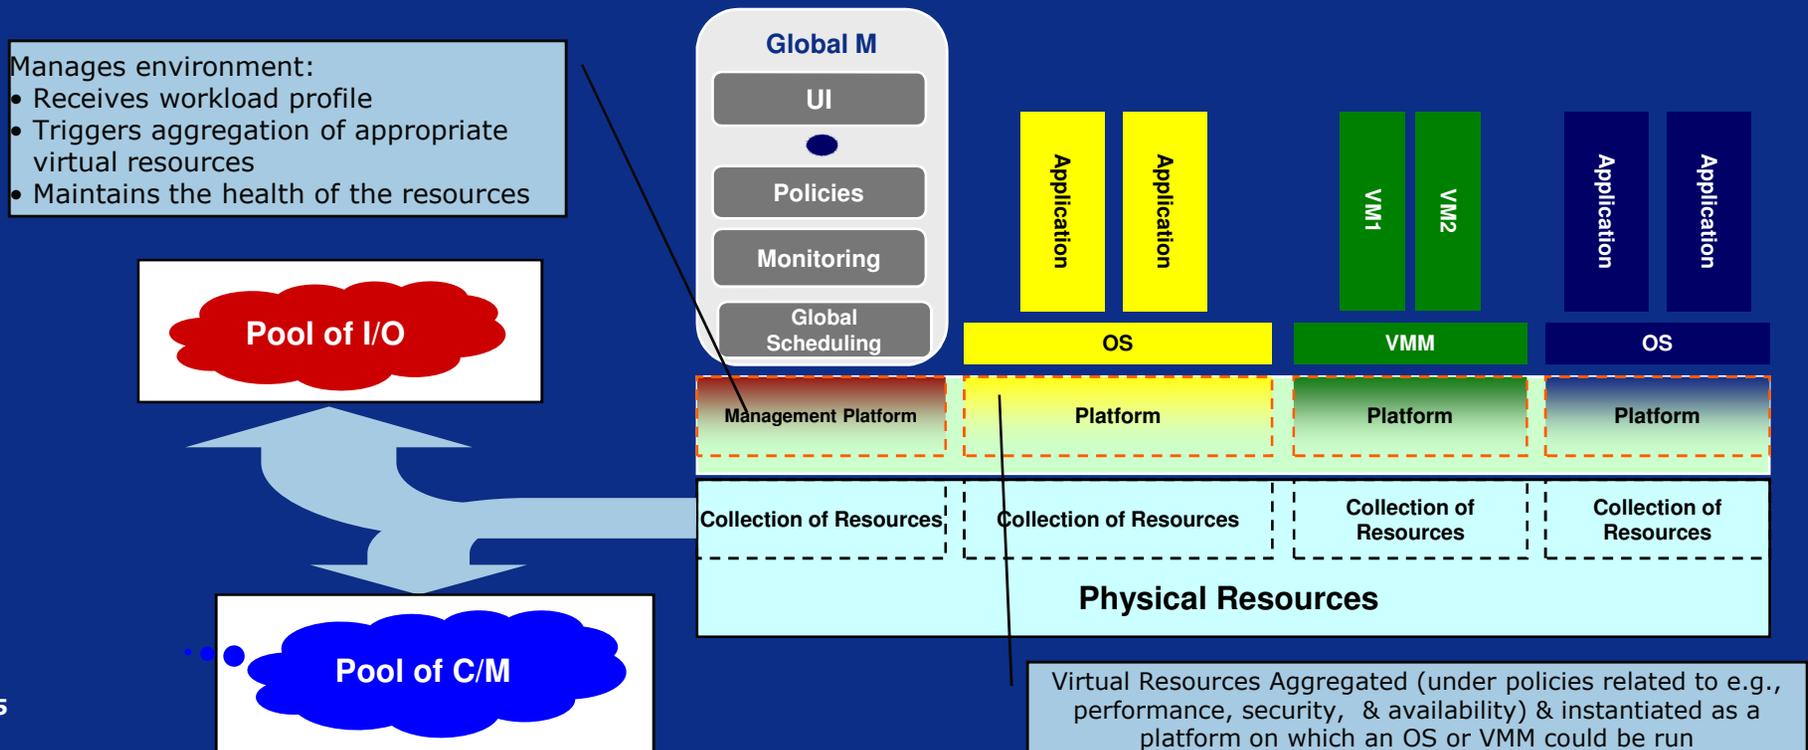
# Autonomic Computing



Note that the petals protrude outside the self-managed disc to indicate that we cannot Self\* all scenarios.

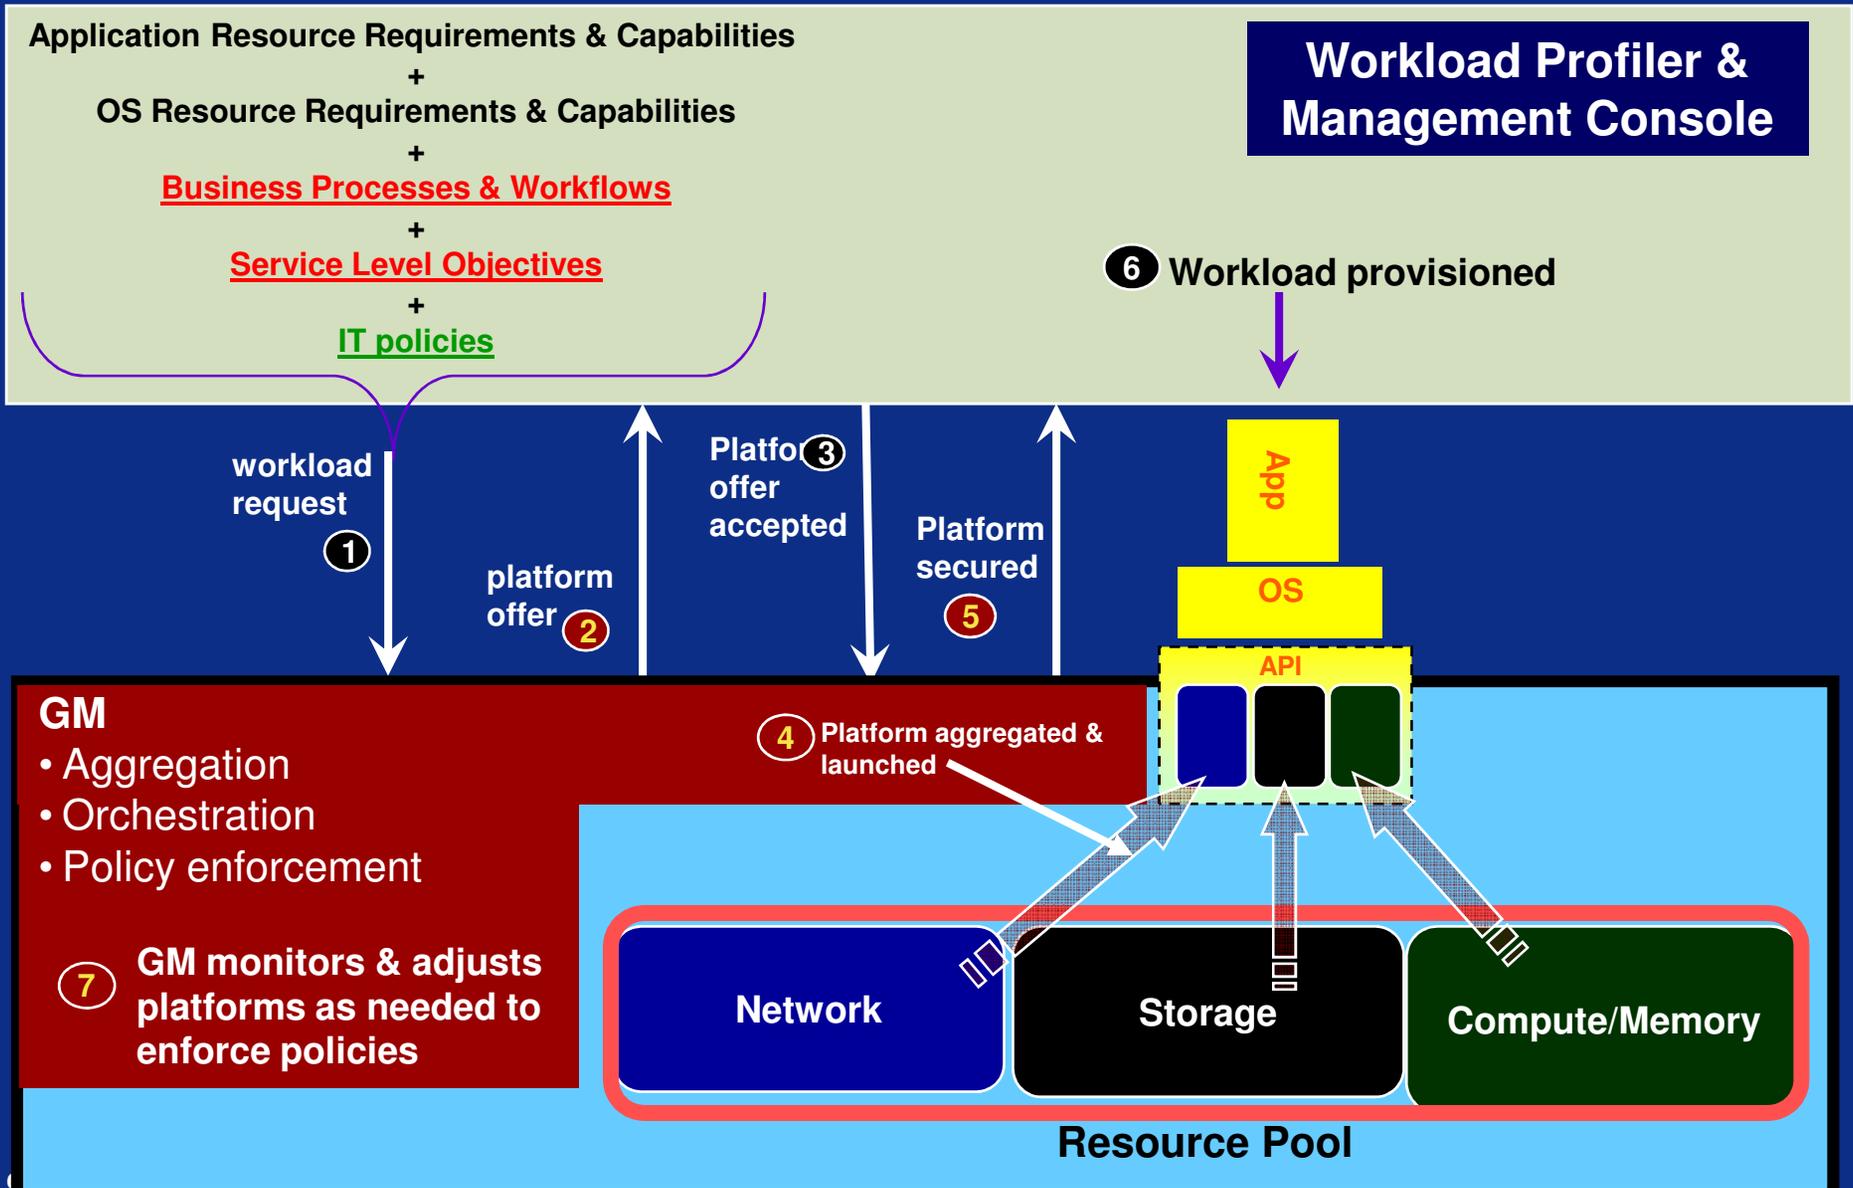
# Future Virtualized Datacenters

- **Dynamic virtual Platforms** from pools of resources
  - Workload management
    - Capacities/features supporting given SLAs, business models & high-level policies based on time, price & performance;
    - Scaling resources up/down to efficiently react to runtime changes (e.g., workload, fault); &
- Managing platforms
  - ... enabling autonomic features & high-level policies
  - Independent of the platforms' system software

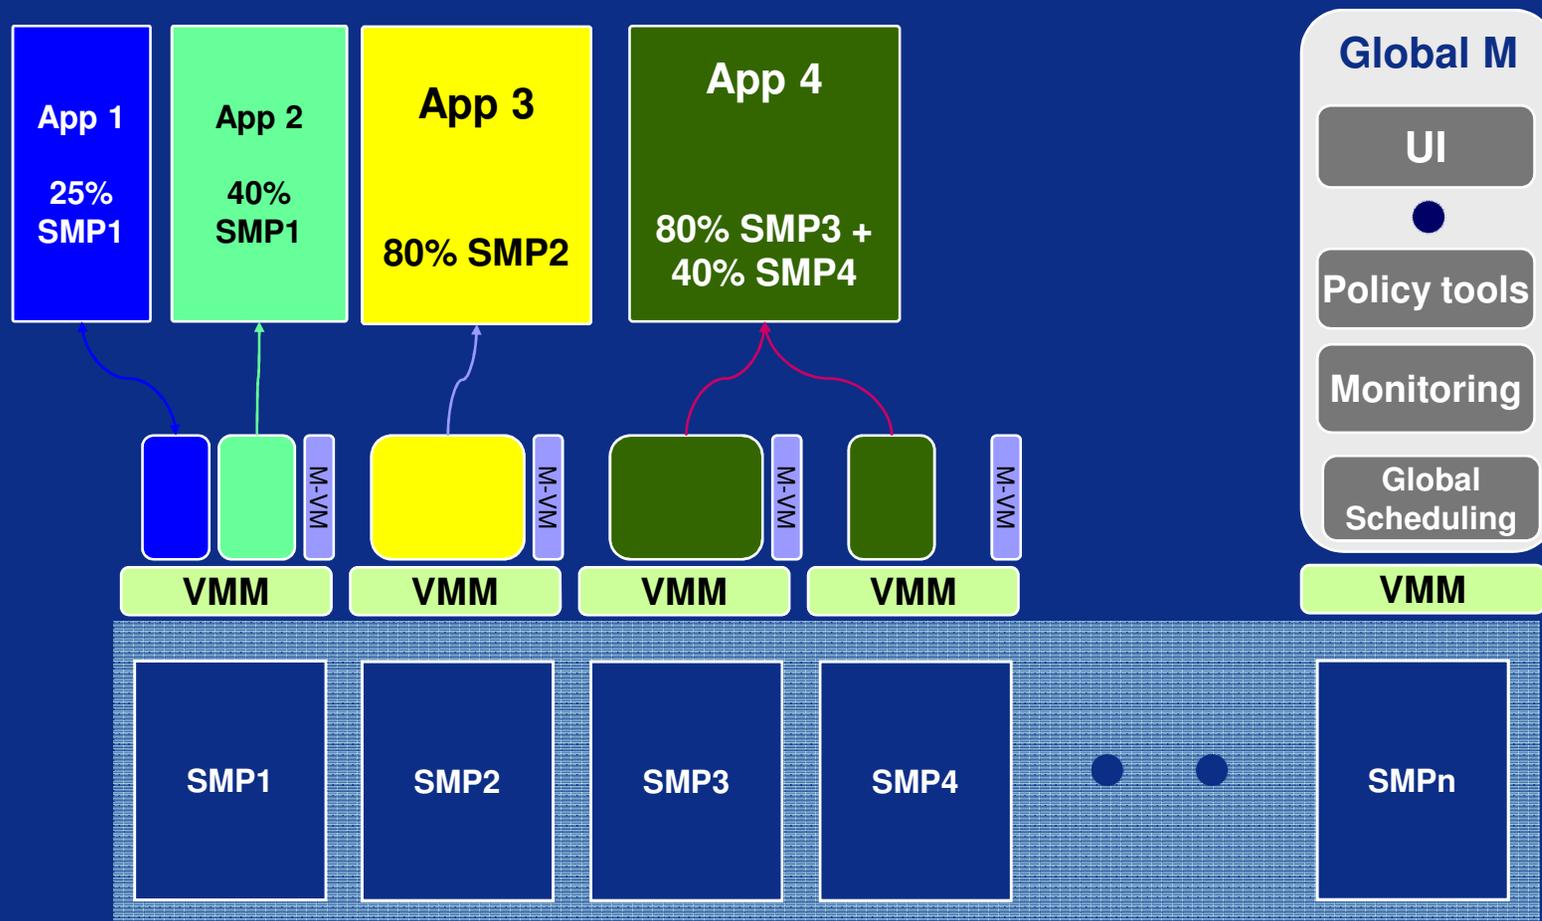


# Future Virtualized Datacenters

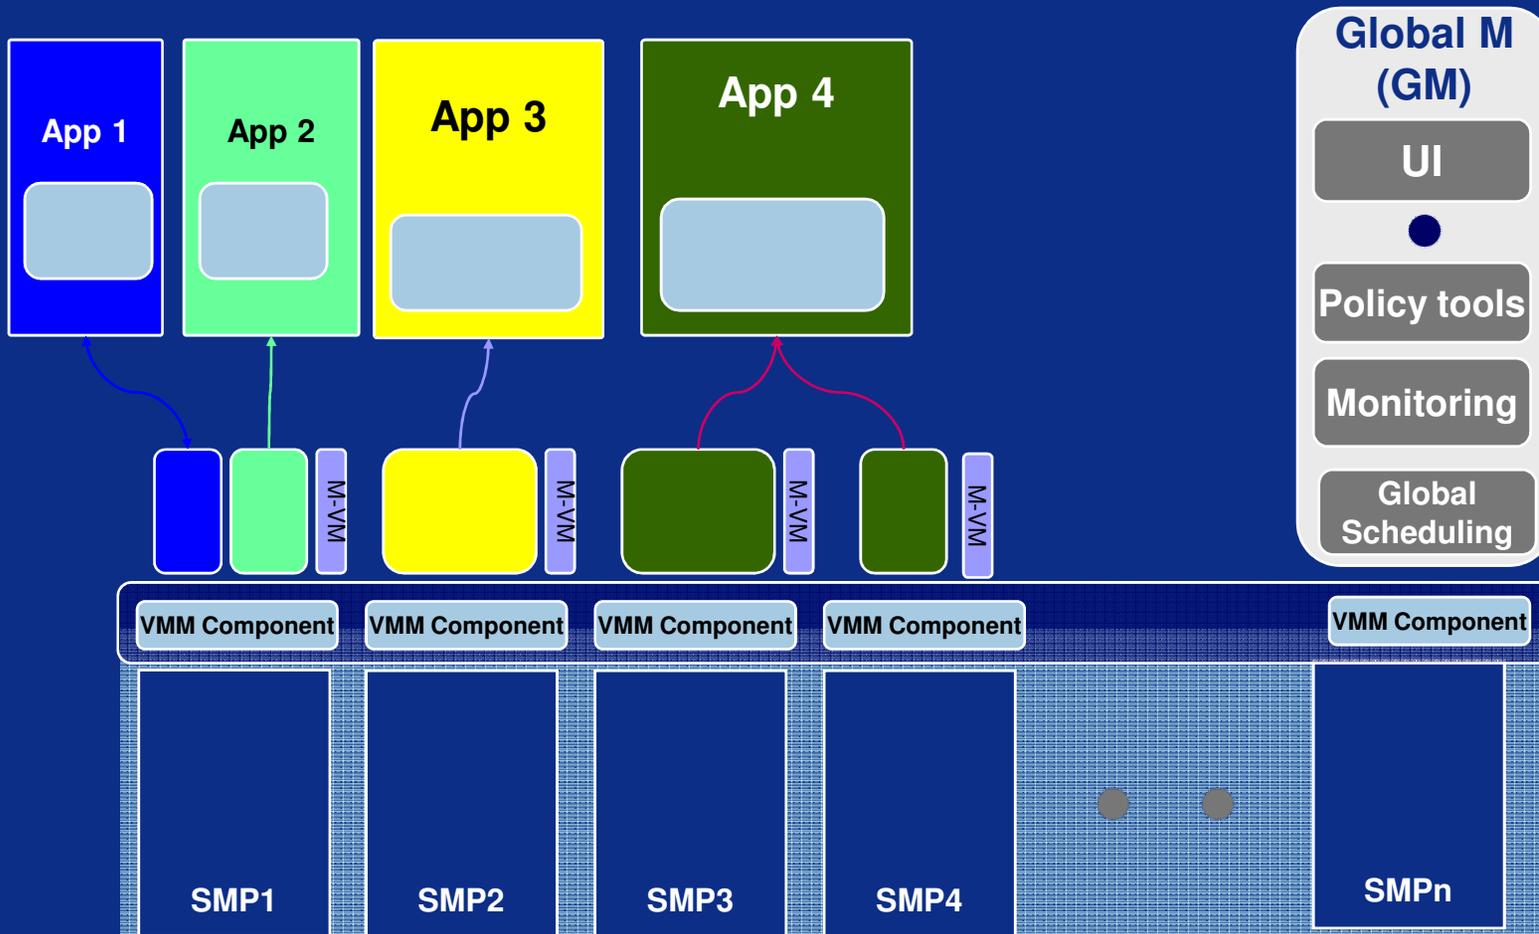
## Example: A Workload Profile



# Future Virtualized Datacenters Architecture I



# Future Virtualized Datacenters Architecture II



## Virtual Enterprise Architecture

### Dynamic Resource Allocation – Sandpiper

---

- **Sandpiper:** automatically detect and mitigate hotspots through virtual machine migration
- *When* to migrate?
- *Where* to move to?
- *How much* of each resource to allocate?
- How much information needed to make decisions?

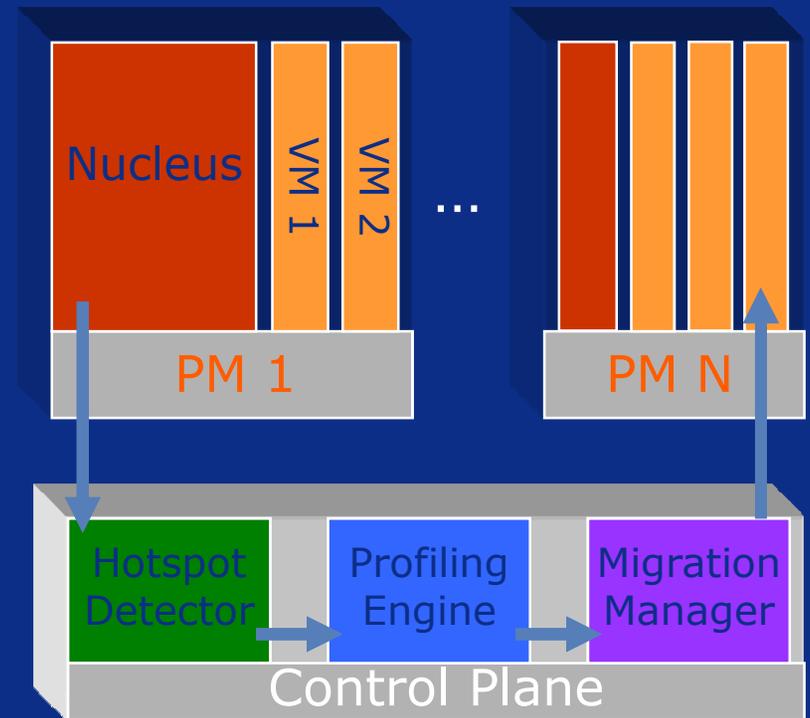


A migratory bird

# Virtual Enterprise Architecture

## Sandpiper Architecture

- Nucleus
  - *Monitor* resources
  - Report to control plane
  - One per server
- Control Plane
  - Centralized server
- Hotspot Detector
  - Detect *when* a hotspot occurs
- Profiling Engine
  - Decide *how much* to allocate
- Migration Manager
  - Determine *where* to migrate



PM = Physical Machine  
VM = Virtual Machine

## Virtual Enterprise Architecture

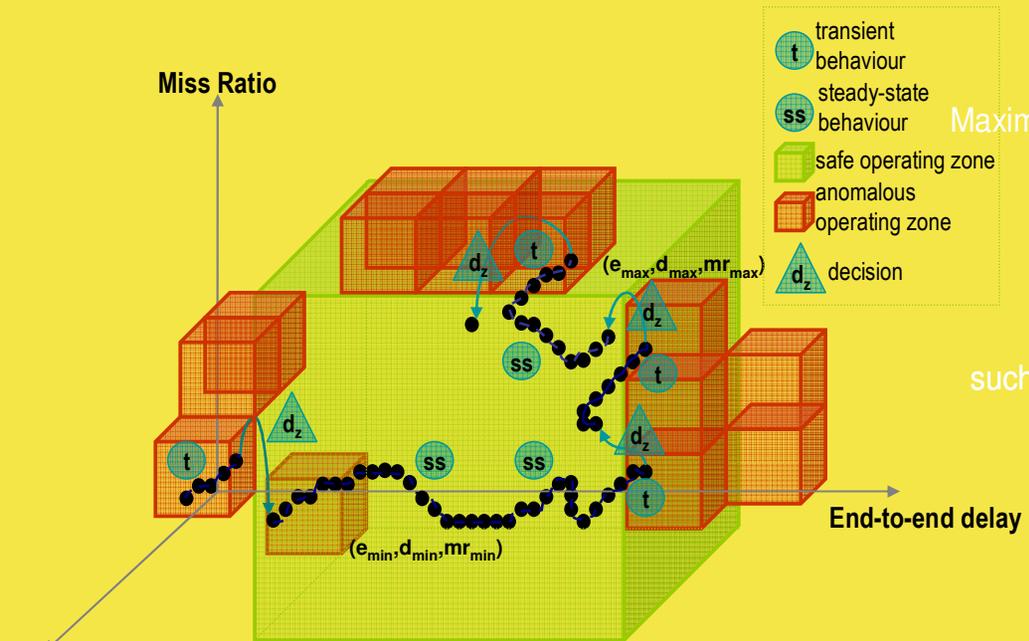
### Sandpiper Black-Box and Gray-Box

- Black-box: only data from outside the VM → Completely OS and application agnostic



- Gray-Box: access to OS stats & application logs
  - Request level data can improve detection and profiling
  - Not always feasible – customer may control OS

# AppFlow: Avirtec Patented Technology to autonomically manage applications/services



- transient behaviour
- steady-state behaviour
- safe operating zone
- anomalous operating zone
- ▲ decision

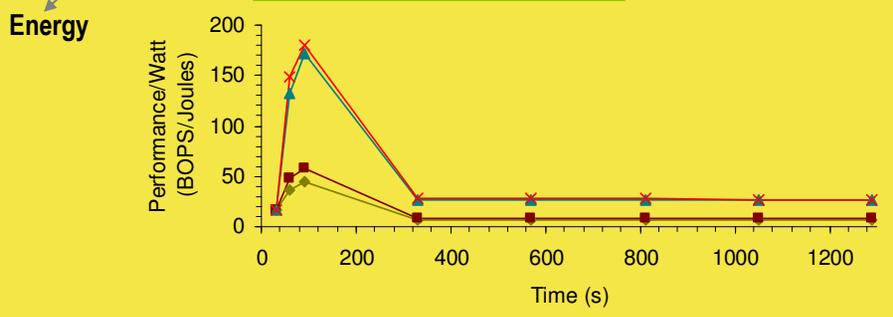
**Maximize performance/watt**

$$ppw_i = \frac{1}{\tau_{a_k} * e_i}$$

$$e_i = \sum_{k=0}^N (c_{jk} * t_{trans_{jk}} + p * n_k * t_{obs}) * X_{jk}$$

such that

1.  $n_k * size / Rank \geq N_{opt} * pageSize$
2.  $Max(\sqrt{\rho_k}) \leq \rho_{th}$   
ch:
3.  $Min(\sqrt{\tau_{a_k}}) \geq \tau_{a_{th}}$
4.  $\sum_{k=0}^n x_{jk} = 1$
5.  $\forall x_{jk} = 0 \vee 1$



◆ PPW (I-16 Ranks)      ■ PPW (II-12 Ranks)  
▲ PPW (III-8 Ranks, 1 Branch)      × PPW (IV-8 Ranks, 2 Branches)

- **Maximum performance/watt improvement of 88.48%**
- **Energy saving of about 48.8 % (26.7 kJ)**

# Summary of Vision

---

*Integrate capabilities & features to datacenters platforms to proactively & efficiently react positively to runtime changes*

- To allow to independently scale resources as workload changes
- To allow to transparently mitigate faults
- To allows resources to self-define and self-configure themselves based on environment

*Our approach lowers TCO: both Capex & Opex*

- Capex through improving the efficiency of resources
- Opex through (Gartner puts Opex at 60-70% of TCO & is primarily labor)
  - Self-management (e.g., self-heal, self-protect, etc.);
  - Through policies such as those reducing energy consumption; &
  - Reducing serviceability and maintainability overhead

# Conclusions

---

- Autonomic Computing is happening in the Enterprise
  - ... However, great deal of both research & technology development challenges still exist
  - Can build the plumbing, but the intelligence will lag
  - Automatic problem recognition - e.g., root cause
  - Policies, Business Processes, SLA to low-level mapping
  - Interoperability will always remain an issue
    - Standards are the right place to work together
    - ... But let's first take a serious look at what we have been doing
- Very strongly advocate collaboration among industry and academia



## NSF Industry/University Cooperative Research (I/UCRC) Program

# The Center for Autonomic Computing

**José A. B. Fortes, UF**

**Salim Hariri, UA**

**Manish Parashar,  
Rutgers**



---

# Q&A

# Thank You

**Mazin Yousif, PhD**  
**Avirtec Corporation**

[mazin@Avirtec.net](mailto:mazin@Avirtec.net)