ATCA System Managers
What, Where, Why

Tutorial T2A
October 2009

Hank Bruning
President
JBlade

System Managers
Why do you need them?

Top 5 System Management Questions

- How much do I have ?
- How much is allocated ?
- How much am I using ?
- Who is using it ?
- Is it broken ?

Questions by Resource

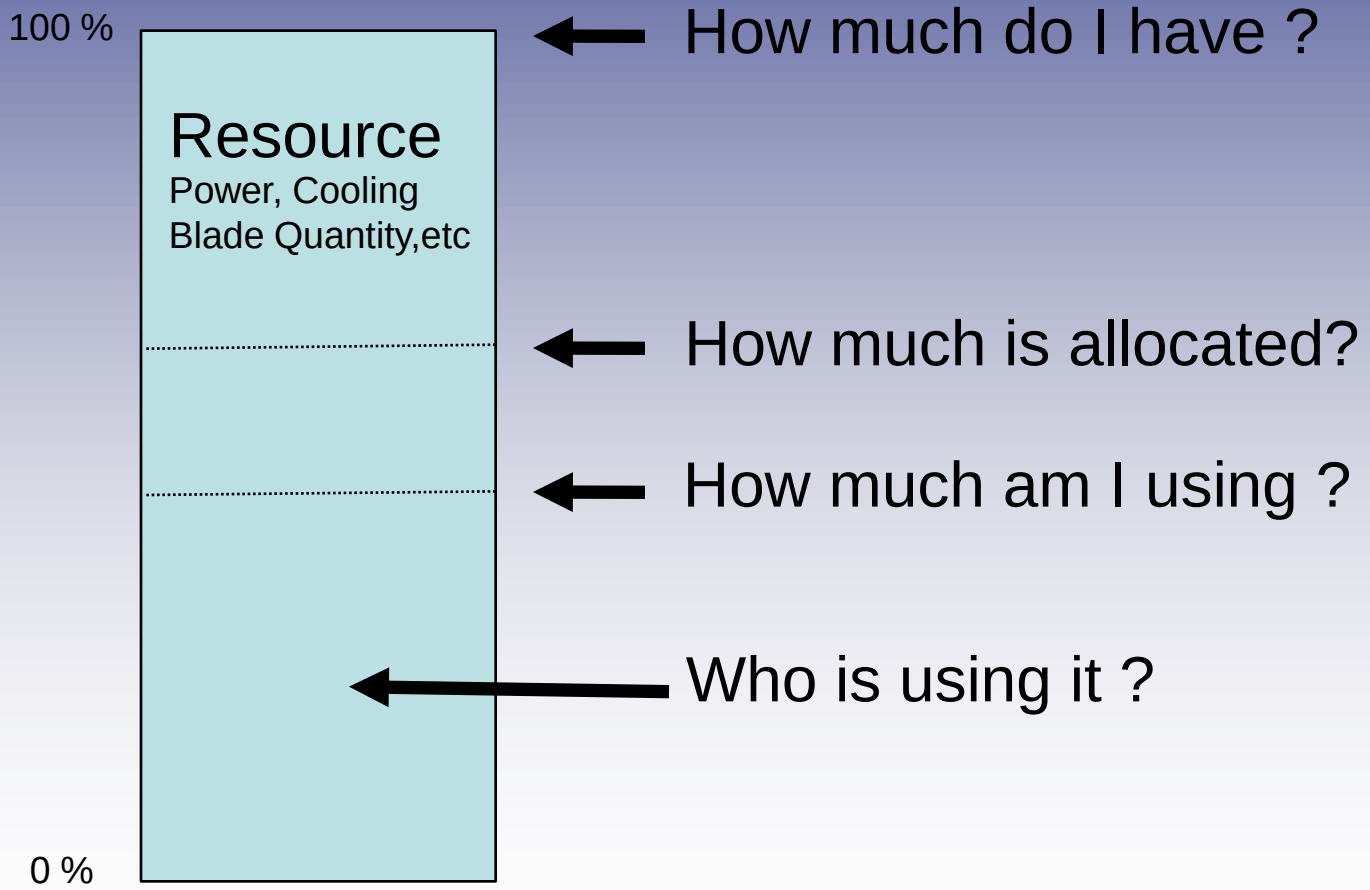

System Manager Feature Matrix

Vendor / Customer Specific

System Manager Feature Matrix

What about Applications ?
More later.....

Question 1:How much do I have ?

Fabric Triple Replicated Mesh

Display Physical Slot 2
Logical Slot 2
Physical Source 0x84

Display Slot Numbers
- ○ as logical Slot Numbers
- ◉ as physical Slot Numbers

Interface Type	Channel Width	Source Physical Slot	Source IPMB	Source Channel	Remote Physical Slot	Remote IPMB	Remote Channel
BASE	FULL	2	0x84	2	1	0x82	2
BASE	FULL	2	0x84	3	3	0x86	2
BASE	FULL	2	0x84	4	4	0x88	2
BASE	FULL	2	0x84	5	5	0x8a	2
BASE	FULL	2	0x84	6	6	0x8c	2
FABRIC	FULL	2	0x84	1	1	0x82	1
FABRIC	FULL	2	0x84	6	1	0x82	6
FABRIC	FULL	2	0x84	11	1	0x82	11
FABRIC	FULL	3	0x86	2	2	0x84	2
FABRIC	FULL	3	0x86	7	2	0x84	7
FABRIC	FULL	3	0x86	12	2	0x84	12
FABRIC	FULL	4	0x88	2	2	0x84	3
FABRIC	FULL	4	0x88	7	2	0x84	8
FABRIC	FULL	4	0x88	12	2	0x84	13
FABRIC	FULL	5	0x8a	2	2	0x84	4
FABRIC	FULL	5	0x8a	7	2	0x84	9
FABRIC	FULL	5	0x8a	12	2	0x84	14
FABRIC	FULL	6	0x8c	2	2	0x84	5
FABRIC	FULL	6	0x8c	7	2	0x84	10
FABRIC	FULL	6	0x8c	12	2	0x84	15
ShMC	FULL	2	0x84	1	Left Shelf M...	0x12	1
ShMC	FULL	2	0x84	2	Right Shelf ...	0x10	2
UPDATE	FULL	1	0x82	1	2	0x84	1

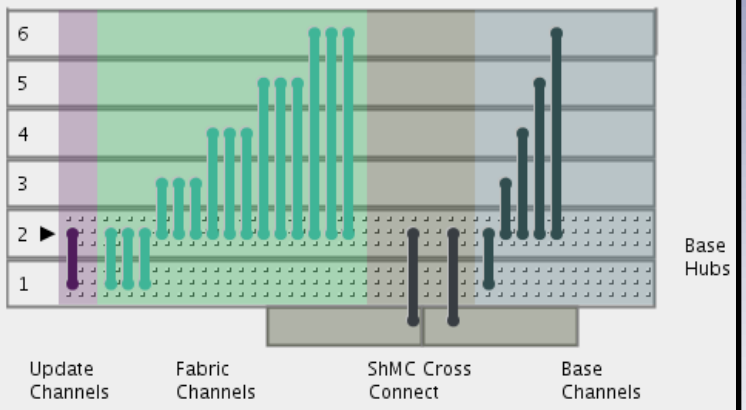

Question 2: How much is allocated ?

Power input to PEM

Power Feed 0

Current power usage	125 Watts
Maximum power usage	1000 Watts
Power utilization	12 Percent
Maximum current available to Shelf	20 Amps
Maximum current Shelf can use	25 Amps
Minimum expected voltage	-40.0 Volts

The installer allocated 20 Amps.

The Shelf is designed for 25 Amps.

Question 3: How much am I using ?

Fans

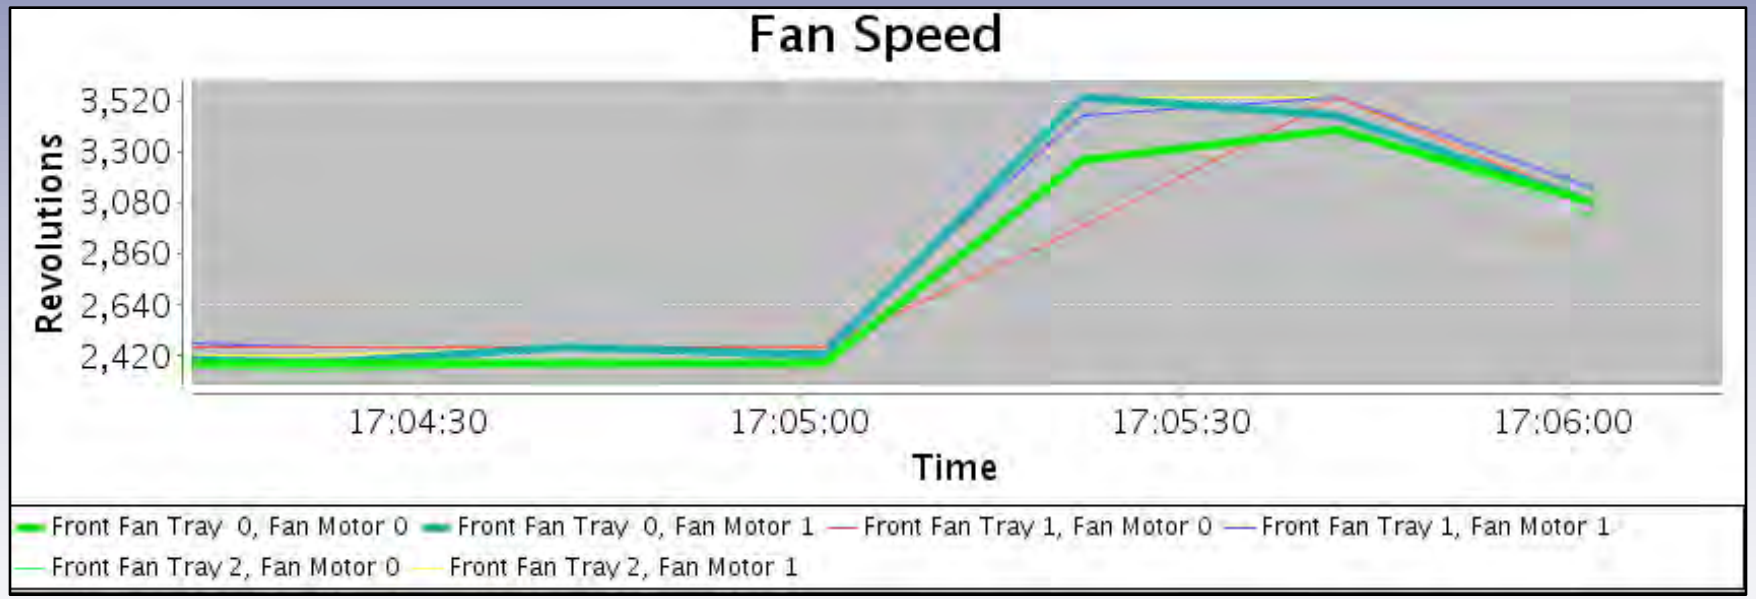

Question 4: Who is using it ?

- FRU Inventory
- Power
- Cooling
- Fabric

Question 5: Is it broken ?

Radial IPMB-0

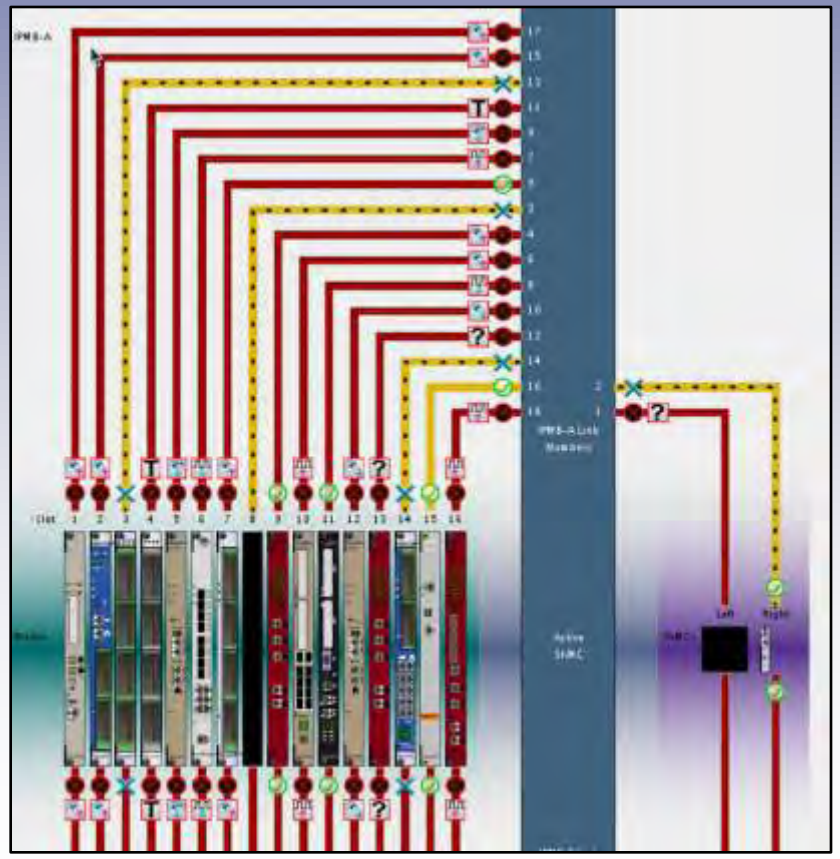

A 16 slot shelf has more than 64 IPMB discrete data points.

Question 5: Is it broken ?

Radial IPMB-0

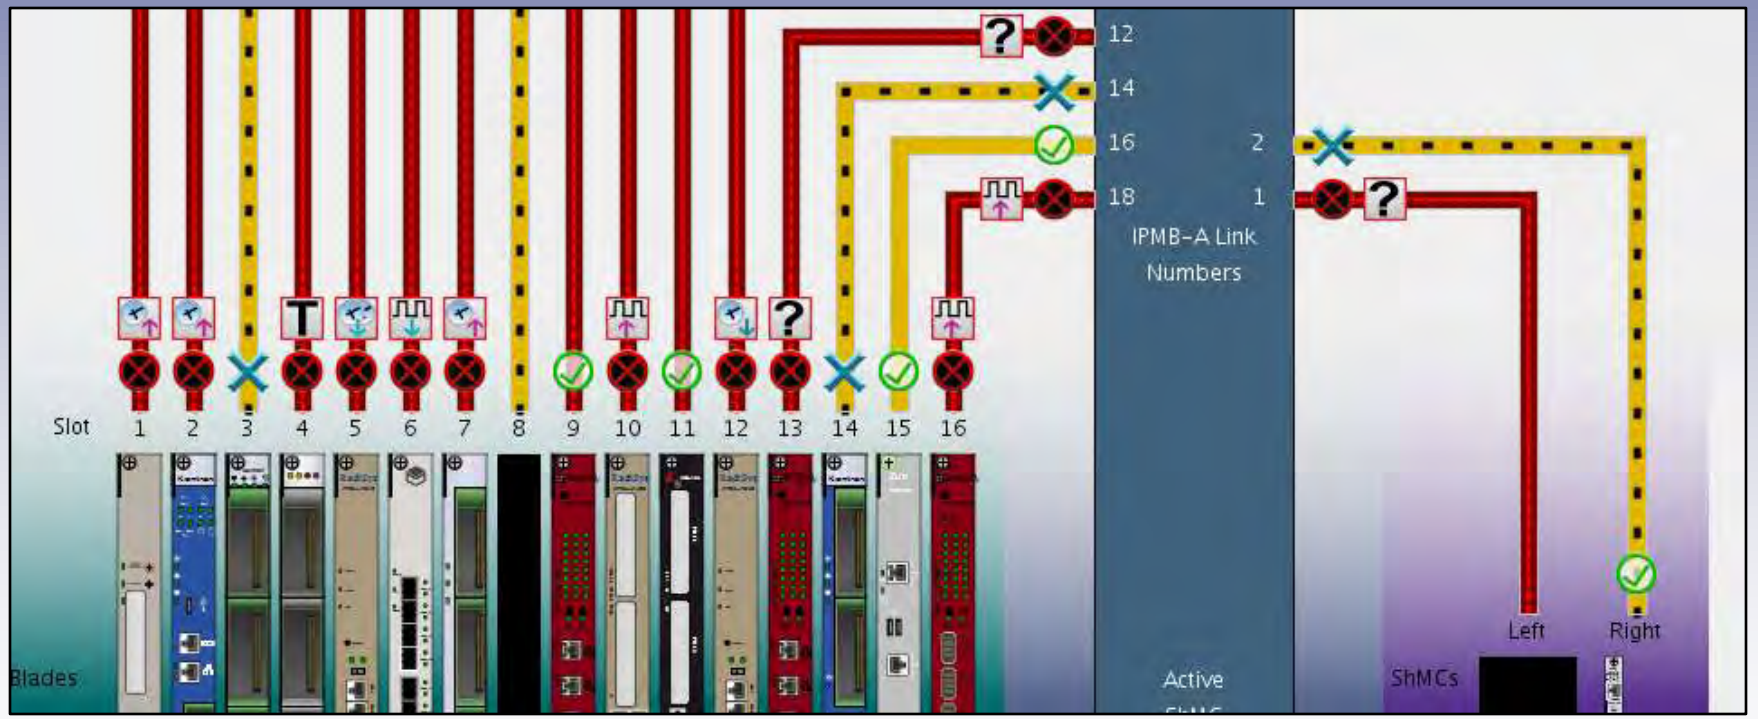

Simulated errors to demonstrate the range of problems.

System Managers
Where do they fit ?

Non Redundant Architecture

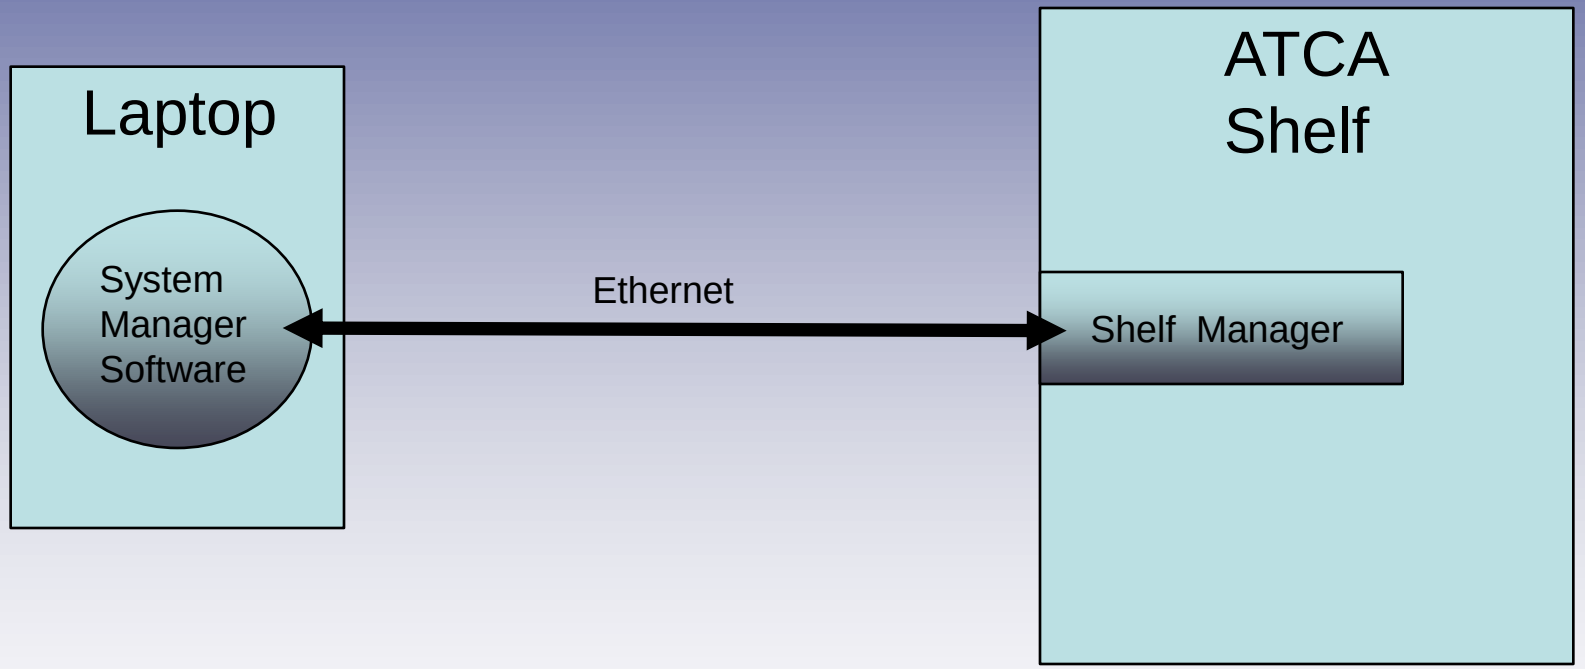

Multi-protocol Architecture

Operator and System Manager are single points of failure.

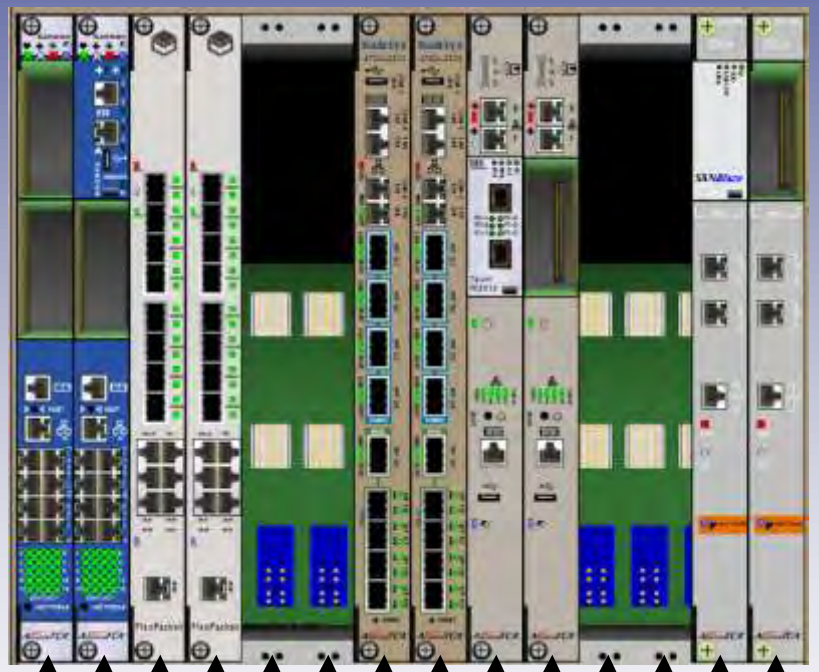

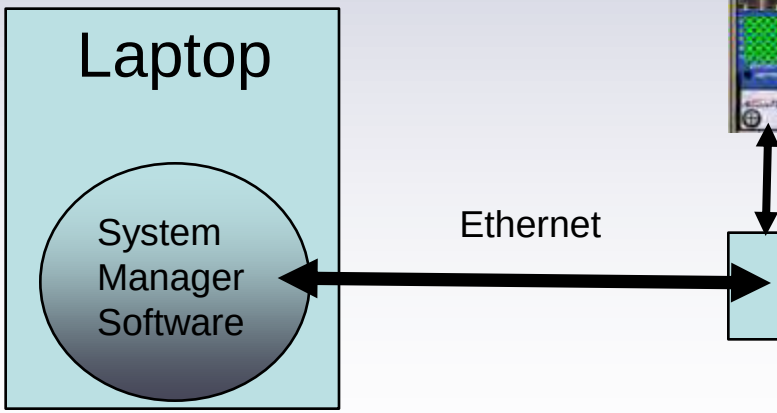

Ethernet

10/100 Ethernet Switch

Redundant Architecture

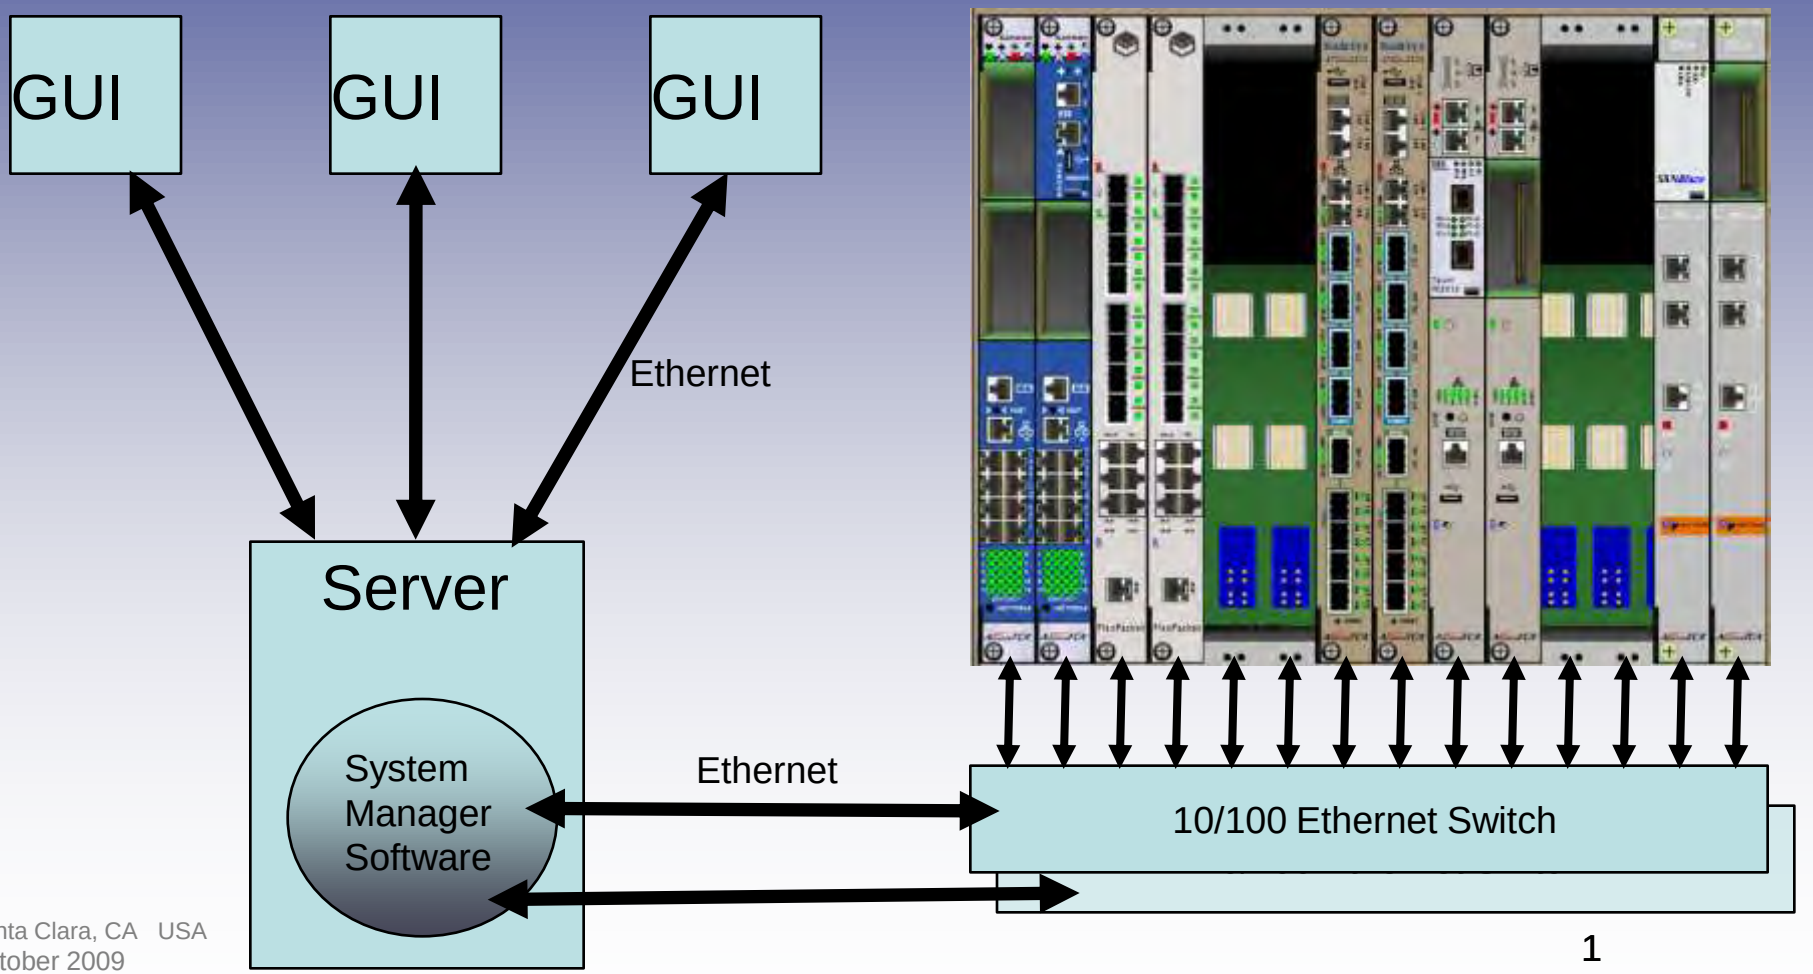

System Managers
What are they?

PICMG 3.0 System Manager Support

- A shelf does not need a System Manager to run
- IPv4 packets containing IPMI RMCP
- PICMG 3.0 provides a mechanism not policy
  1) Fan control
  2) Multi protocol EKeying
  3) FRU activation (save power by not activating, big system power on)
  4) Shelf Manager power allocation

Application Support

- System Managers application support
  1) No applications
  2) SAF
  3) VMWare

- Number of ATCA Blades supporting VMWare
  October 2008, zero
  October 2009, four

Systemic problem: ATCA Platform management provides data not available via any application aware API.

System Manager Generations

	RMCP	RMCP SNMP	RMCP, SNMP, NetConf, SAF, VMi, NetFlow
	Don't care. Any communication is good	Messages/Second	Messages/Millisecond
	Shelf Manager	Shelf Manager	Shelf Manager & Operator
	IPMB-0 via Shelf Manager	IPMB-0, Base Fabric	Unique cable per blade
	IPMI Thresholds	IPMI Thresholds and discrete	IPMI Thresholds and discrete. ATCA state and fabric use

Third Generation System Managers

Required

- Three or more protocols. RMCP,SNMP, NetConf, SAF, VMi, NetFlow
- Peek RMCP speed 3 packets/millisecond
- Multiple Ethernet connections to each switch
- Graphs of analog sensors

Optional

- Redundancy
- Multiple Ethernet connections to all blades
- HPM.1 Firmware upgrades
- Multi user

Third Generation Bandwidth Usage

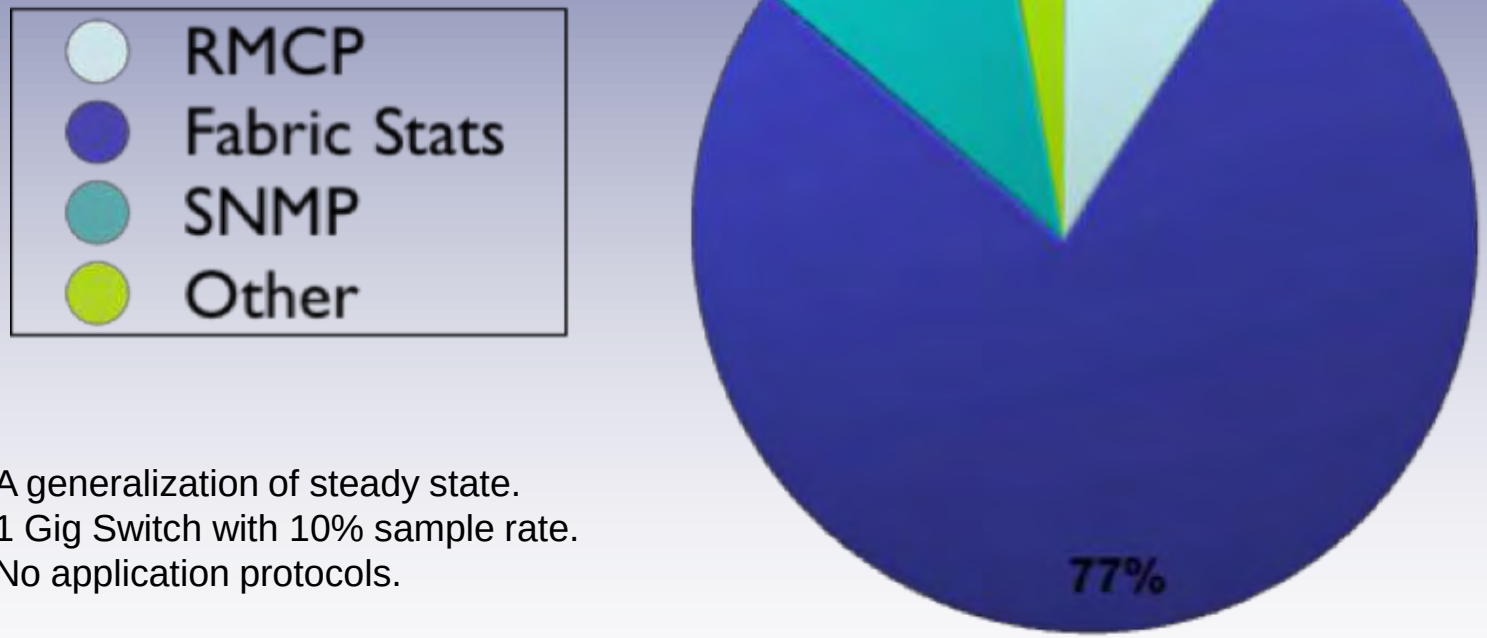

A generalization of steady state.
1 Gig Switch with 10% sample rate.
No application protocols.

System Manager Review

- They provide information on capacity, allocation and usage of resources in the shelf.
- They reside outside the shelf and cooperate with data collectors inside the shelf.
- The industry is now developing their third generation system managers.

Additional Information

- ATCA Hardware supporting VMWare
  http://www.vmware.com/resources/compatibility/get_pdf.php?deviceCategory=server
- Service Availability™
  http://www.saforum.org
- NetConf
  http://www.ietf.org/rfc/rfc4741.txt

Thank you.

Hank Bruning

hank.bruning@jblade.com