Monitoring z/VM with SNMP

Session 8448

Jay Brenneman
rjbrenn@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

DB2*
DB2 Connect
DB2 Universal Database
e-business logo*
e-business on demand
HiperSockets
IBM*
IBM eServer
IBM logo*
IMS
Resource Link
S/390*
Tivoli*
Tivoli Storage Manager
TotalStorage*
WebSphere*
z/OS*
z/VM*
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
UNIX is a registered trademark of The Open Group in the United States and other countries.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- SNMP and you
- z/VM SNMP capabilities
- z/VM Setup details
- SNMP Network Monitoring Stations

SNMP and you

- Simple Network Monitor Protocol
  - Identifies the host and provides admin contact information
  - Provides port numbers and addresses of network interfaces
  - Sends notifications of link state changes
- The standard network monitor used by network admins everywhere. Good for getting buy in from network team, it makes VM less foreign.
- A good fit for VM: sends notifications when it needs to and uses no CPU when sleeping
- Uses a tree data structure:

The SNMP Tree

A tree of OIDs (Object Identifier)
Dot notation:
.1.3.6.1.2.1.1.1
ASN notation:
.iso.org.dod.internet.mgmt.mib-2.system.sysDescr

The sysDescr of one of my VM Systems:
IBM 2097
z/VM Version 6 Release 1.0, service level 0901 (64-bit)
VM TCP/IP Level 610
RSU 0901 running TCPIP MODULE M2,
dated 10/29/09 at 16:47

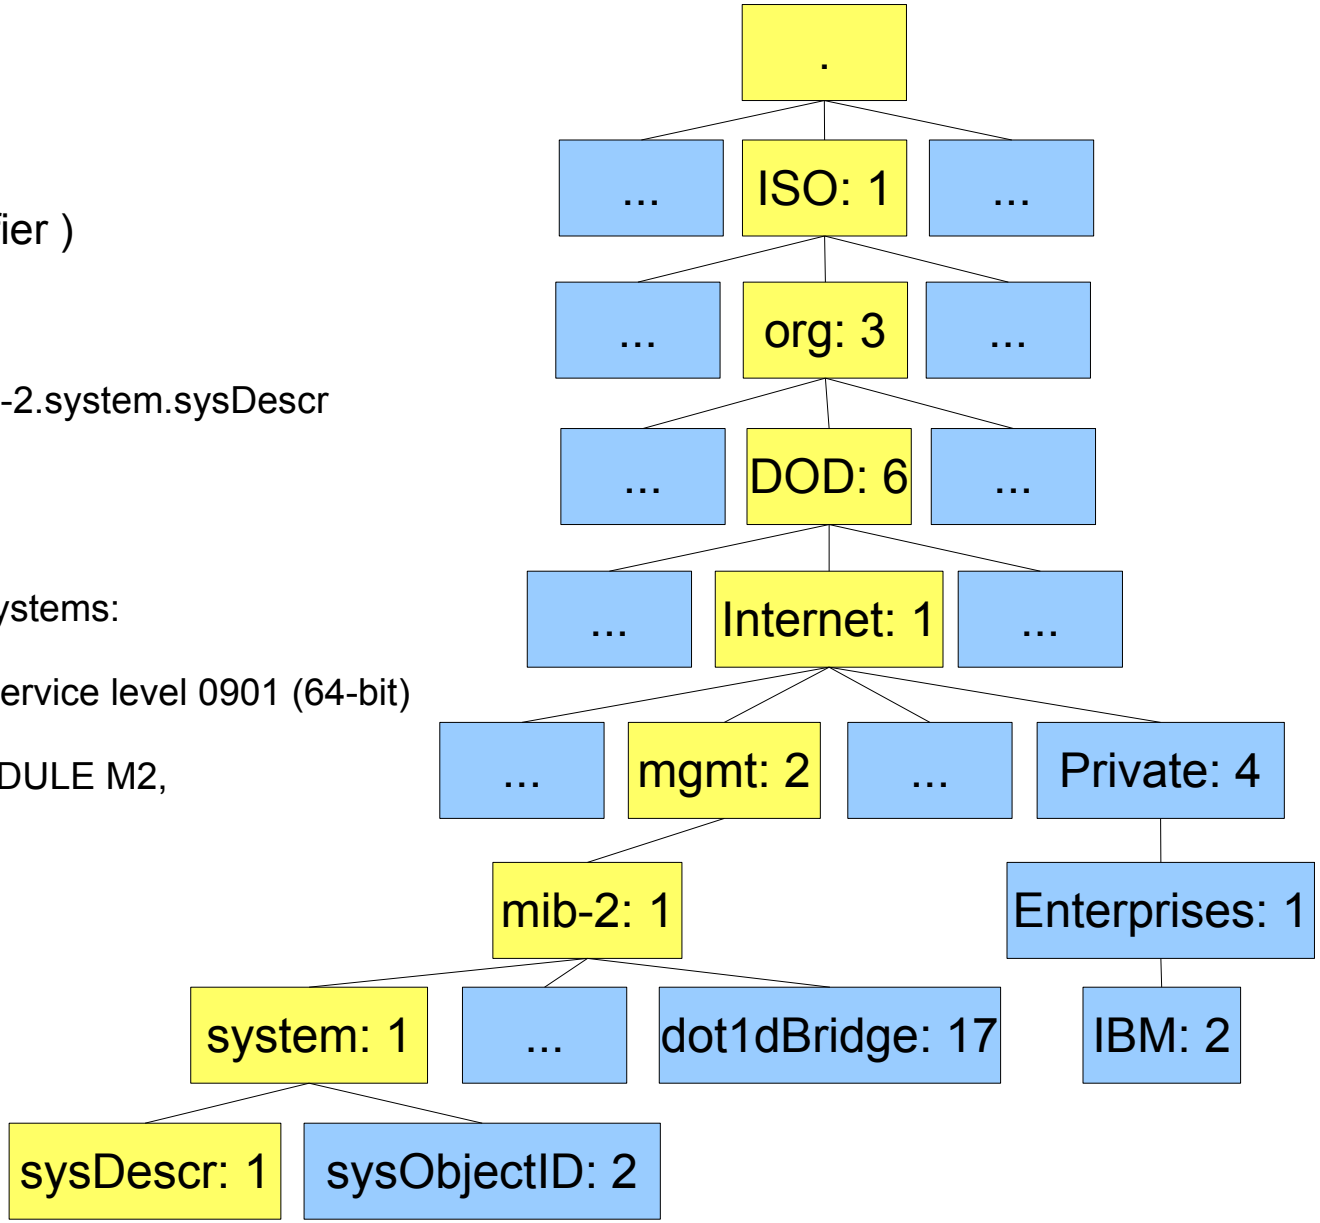

z/VM's SNMP capabilities

- z/VM's SNMP agent is already used by existing management tools
- Link State Changes (traps)
  - Sep 1 11:13:26 litnetm1 snmptrapd[1178]: 2009-09-01 11:13:26 192.168.70.28(via UDP: [192.168.71.15]:161) TRAP, SNMP v1, community LTICVM9 IF-MIB::ifIndex Link Down Trap (0) Uptime: 5 days, 15:36:55.00 IF-MIB::ifIndex = INTEGER: 1
- mib-2 data provided:
  - System: sysDescr, sysObjectID, sysUptime, sysContact, sysName, sysLocation
  - Interfaces: ifNumber, ifTable: ifIndex, ifType, ifDescr, ifMtu, ifSpeed, ifPhysAddress, ifAdminStatus, ifOperStatus, ifLastChange, ifInOctets, ifInUcastPkts, ifInDiscards...
  - IP: ipForwarding, ipInReceives, ipInDiscards, ipOutRequests, ipAddrTable, ipRouteTable...
  - ICMP: icmpInMsgs, icmpOutMsgs, icmpInEchos, icmpOutEchos, icmpOutDestUnreachs...
  - TCP: tcpActiveOpens, tcpPassiveOpens, tcpRetranSegs, tcpInSegs, tcpOutSegs, tcpConnTable...
  - UDP: udpInDatagrams, udpOutDatagrams, udpTable...

z/VM's SNMP capabilities

- z/VM 5.4 added support for a SNMP Subagent which provides access to VSWITCH counters.
  - APAR VM64646 required to fix CP's responses to the subagent so that it does not abend
- dot1dBridge data provided:
  - dot1dBase: dot1dBasePortAddress, dot1dBaseNumPorts, dot1dBasePortTable, dot1dBasePortDelayExceededDiscards, dot1dBasePortMTUExceededDiscards...
  - dot1dTp: dot1dTpPortTable: dot1dTpPortMaxInfo, **dot1dTpPortInFrames, dot1dTpPortOutFrames, dot1dTpPortInDiscards**
  - Most useful is likely to be the port counts
    They provide the same data as seen in Q VSWITCH <vswname> USER <userid>

z/VM's SNMP limitations

- SNMP v1 only
  - Plain text community names in all requests / responses
  - Data structure limitations in v1 tree prevent some network configurations from being “walkable”
- Currently no way to tie a Vswitch port number to a guest
  - Guests couple in to the lowest numbered port starting at port 65
  - Data returned by Q VSWITCH is stale as soon as its returned: a point in time snapshot
  - Some SNMP monitor stations may be able to connect the dots based on MAC addresses
  - Port count values are read only, and cannot be reset.
    Uncouple / Couple zeroes these values, but also changes the port number that the NIC is plugged in to.
  - Ports are ephemeral – they only exist when a guest NIC is coupled to them.

z/VM SNMP setup details

- Apply or verify APAR VM64646
  - CP change: requires an IPL
- Requires one OSA device for each Vswitch to be monitored
  - This should be through a different port than the one which services the Vswitch!
- TCPIP Profile changes:
  - Add a Device & Link for the monitoring OSA if you're not using an existing link
  - Add SYSCONTACT and SYSLOCATION statements to fill in the system mib data
  - Add SNMPD to the AUTOLOG section
  - Add SNMPD to the PORT section
  - Add SNMPD to the OBEY section
  - Add the VSWITCH statement to the HOME definition for the LINK
  - Start the new Device

z/VM SNMP setup details

- Add the following to SYSTEM DTCPARMS:

```
:nick.SNMPD :type.SERVER :class.snmp
  :owner.TCPMAINT
  :parms.-s SNMPSUBA
:nick.SNMPSUBA :type.SERVER :class.snmp_agent
  :owner.TCPMAINT
  :parms.-u SNMPD
```

- Create PW SRC file accessible by SNMPSUBA:

*Community	*network	*netmask
TICLNET	192.168.71.249	255.255.255.255
TICLNET	192.168.71.48	255.255.255.255
TICLNET	192.168.71.49	255.255.255.255

- Create SNMPTRAP DEST file accessible by SNMPSUBA:

*Host	*Protocol
192.168.71.249	UDP
192.168.71.48	UDP
192.168.71.49	UDP

*check the manuals
You may not be able to use comments like this

z/VM SNMP setup details

- Copy MIB_EXIT SDATA from TCPMAINT's 591 disk to MIB_EXIT DATA on TCPMAINT's 198 disk
- Copy MIB_DESC SDATA from TCPMAINT's 591 disk to MIB_DESC DATA on TCPMAINT's 198 disk
- Copy MIBX2DSC SAMPEXEC from TCPMAINT's 592 disk to MIBX2DSC EXEC on TCPMAINT's 592 disk
- Run MIBX2DSC to copy some of the statements from MIB_EXIT DATA to MIB_DESC DATA
  - MIBX2DSC MIB_EXIT DATA L MIB_DESC DATA L
- If you did it right there will be a new section at the bottom of MIB_DESC DATA containing the bridge OIDs: 1.3.6.1.2.1.17

z/VM SNMP setup details

- Verify that SNMPSUBA has class E privleges
- FORCE and XAUTOLOG SNMPD to verify that it brings up SNMPSUBA automatically
- Run snmpwalk from one of the PW SRC network monitor stations
  - snmpwalk -t 10 -c TICLNET -v 1 192.168.70.24 .1.3.6.1.2.1.1
  - snmpwalk -t 10 -c TICLNET -v 1 192.168.70.24 .1.3.6.1.2.1.17

```
SNMPv2-MIB::sysDescr.0 = STRING: IBM 2097; z/VM Version 5
Release 4.0, service level 0901 (64-bit), VM TCP/IP Level 540; RSU
0901 running TCPIP MODULE M2 dated 05/28/09 at 11:41
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-
SMI::enterprises.2.2.1.2.3
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (180428500)
20 days, 21:11:25.00
SNMPv2-MIB::sysContact.0 = STRING:  BOB ADMIN (TL 555-1122)
GARY SYSPROG (TL 555-1133) DEPARTMENT OF REDUNDANCY
DEPARTMENT
SNMPv2-MIB::sysName.0 = STRING: LTICVM9.PDL.POK.IBM.COM
SNMPv2-MIB::sysLocation.0 = STRING:  BUILDINGA 123 FAKE
STREET POUGHKEEPSIE, NY 12601
SNMPv2-MIB::sysServices.0 = INTEGER: 76
```

```
BRIDGE-MIB::dot1dBaseBridgeAddress.0 = Hex-STRING: 02 09 00 00 00 03
BRIDGE-MIB::dot1dBaseNumPorts.0 = INTEGER: 24
BRIDGE-MIB::dot1dBaseType.0 = INTEGER: transparent-only(2)
BRIDGE-MIB::dot1dTpPort.1 = INTEGER: 1
BRIDGE-MIB::dot1dTpPort.65 = INTEGER: 65
BRIDGE-MIB::dot1dTpPort.66 = INTEGER: 66
BRIDGE-MIB::dot1dTpPort.67 = INTEGER: 67
BRIDGE-MIB::dot1dTpPortMaxInfo.1 = INTEGER: 9152
BRIDGE-MIB::dot1dTpPortMaxInfo.65 = INTEGER: 65472
BRIDGE-MIB::dot1dTpPortMaxInfo.66 = INTEGER: 65472
BRIDGE-MIB::dot1dTpPortMaxInfo.67 = INTEGER: 65472
BRIDGE-MIB::dot1dTpPortInFrames.1 = Counter32: 180755120
BRIDGE-MIB::dot1dTpPortInFrames.65 = Counter32: 17395512
BRIDGE-MIB::dot1dTpPortInFrames.66 = Counter32: 17238306
BRIDGE-MIB::dot1dTpPortInFrames.67 = Counter32: 22867486
BRIDGE-MIB::dot1dTpPortOutFrames.1 = Counter32: 228351243
BRIDGE-MIB::dot1dTpPortOutFrames.65 = Counter32: 6487738
BRIDGE-MIB::dot1dTpPortOutFrames.66 = Counter32: 2554372
BRIDGE-MIB::dot1dTpPortOutFrames.67 = Counter32: 13075687
```

PROFILE TCPIP (partial) example

```
SYSCONTACT
  Bob Admin   (TL 555-1122)
  Gary Sysprog (TL 555-1133)
  Department of Redundancy Department
ENDSYSCONTACT

SYSLOCATION
  BUILDINGA Floor 2 Red 11 Blue 48
  123 Fake Street
  Poughkeepsie, NY 12601
ENDSYSLOCATION

AUTOLOG ( trimmed )
  SNMPD     password    ; SNMP VM Agent Virtual Machine

PORT ( trimmed )
  161  UDP SNMPD              ; SNMP Agent

OBEY
   SNMPD
ENDOBEY

  DEVICE  DEVETH4    OSD      0800
  LINK    LNKETH4    QDIOETHERNET  DEVETH4

HOME
  192.168.70.24                VSWITCH DT70TAG  LNKETH4

  START DEVETH4
```

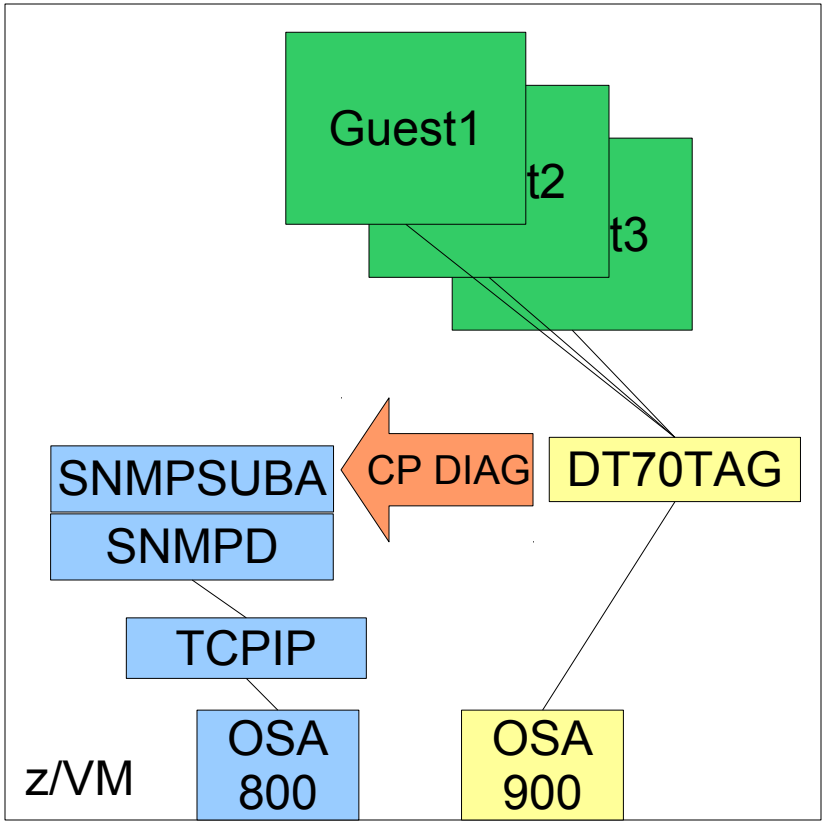

PROFILE TCPIP (partial) example: more vswitches!

```
DEVICE  DEVETH4    OSD     0800
LINK    LNKETH4    QDIOETHERNET  DEVETH4

DEVICE  DEVETH5    OSD     0804
LINK    LNKETH5    QDIOETHERNET  DEVETH5

DEVICE  DEVETH6    OSD     0808
LINK    LNKETH6    QDIOETHERNET  DEVETH6

HOME
  192.168.70.24              VSWITCH DT70TAG  LNKETH4
  192.168.70.25              VSWITCH PRVV68   LNKETH5
  192.168.70.26              VSWITCH PRVV71L2 LNKETH6

  START DEVETH4
  START DEVETH5
  START DEVETH6
```

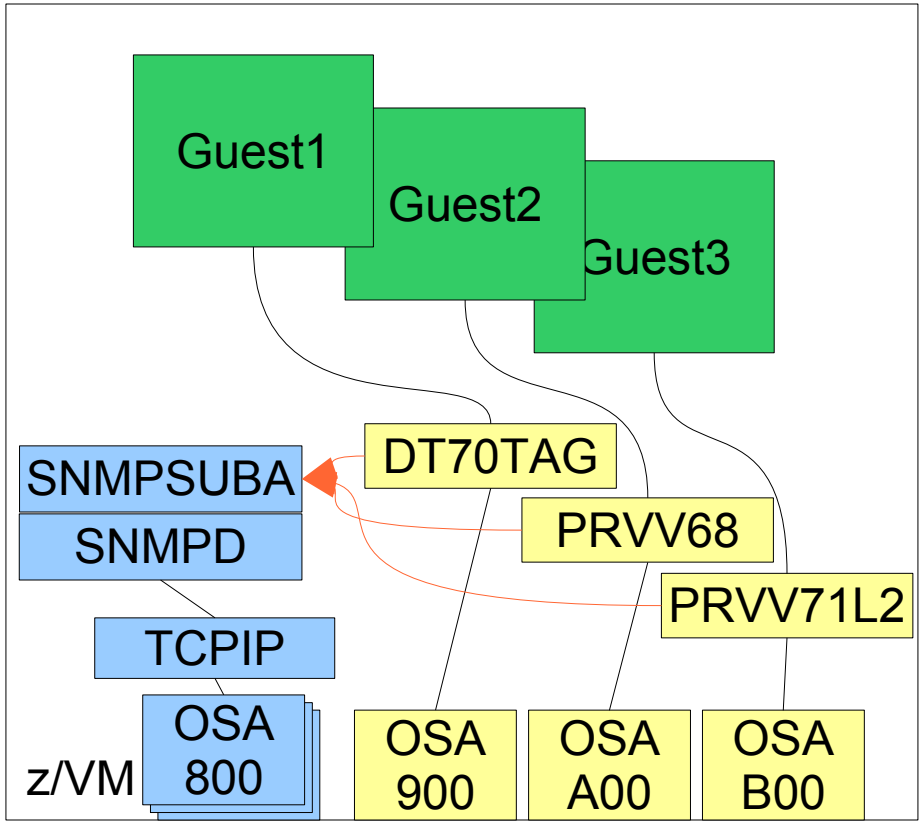

- This appears as a single device to the SNMP Monitor station!
- Security implications if the three Vswitches are in different security zones?

Example: more vswitches and more security!

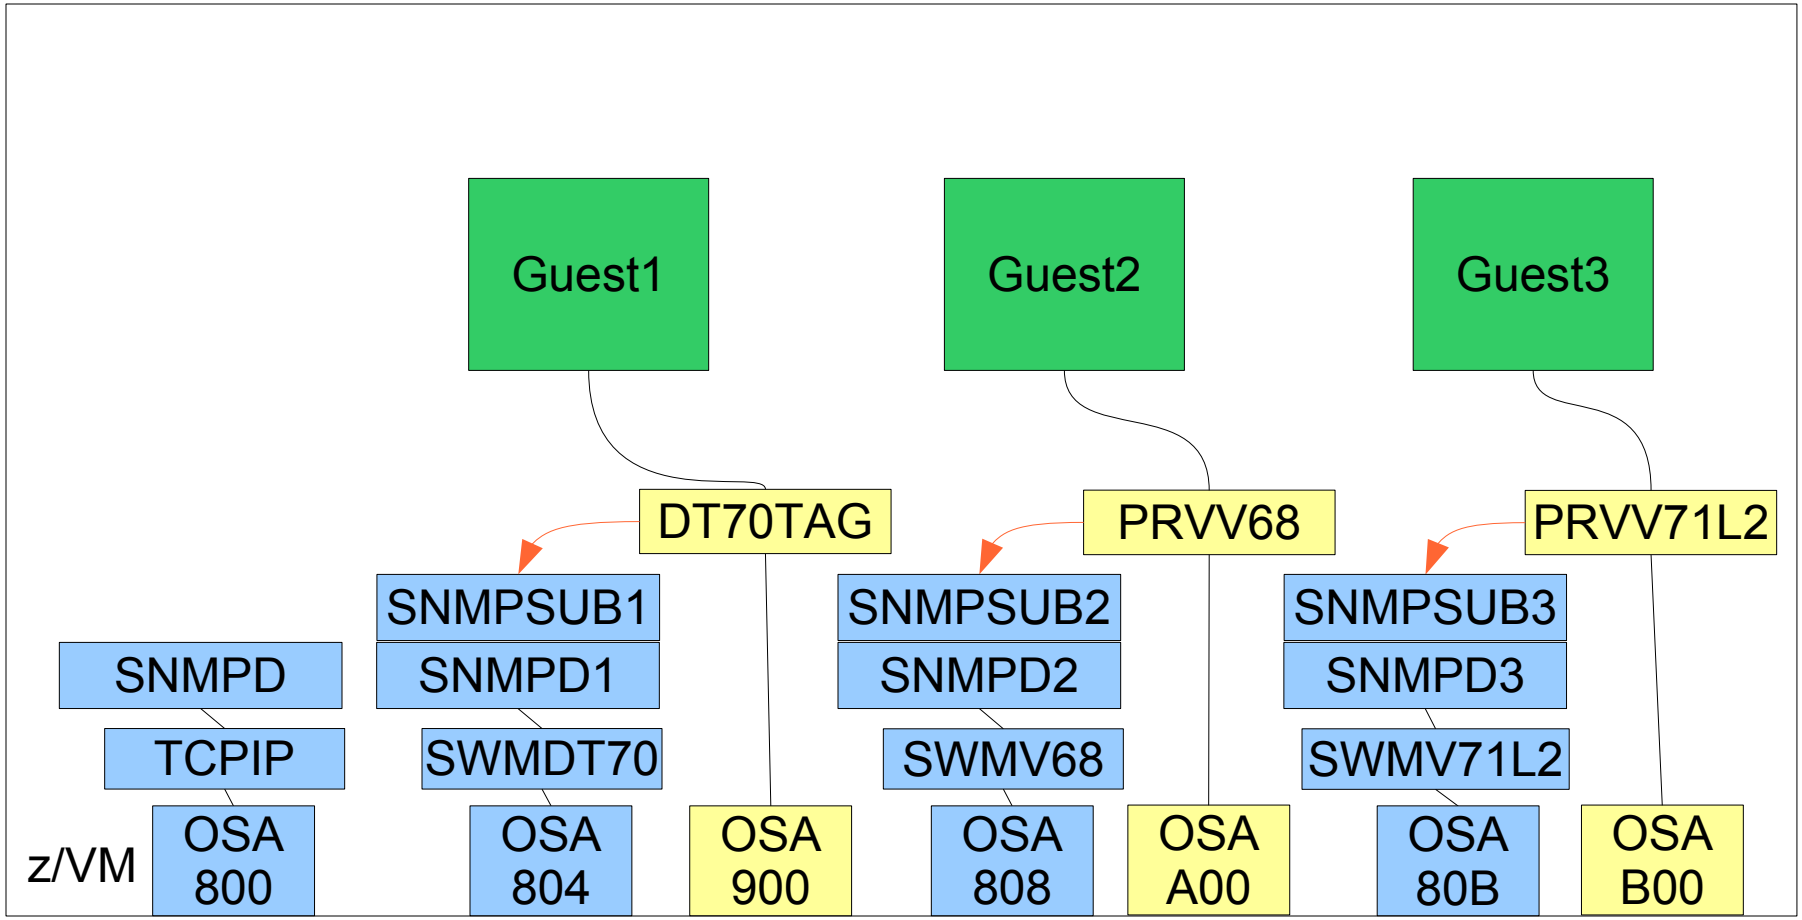

SNMP Network Monitor Stations

- Snmptrapd & (Nagios | Xymon)
- OpenNMS

- NetView
- IBM Tivoli Network Manager IP Edition
  - formerly Netcool Precision IP

SNMPtrapd and Nagios / Xymon

- Xymon is the new name for the Hobbit Monitor
- Both Nagios and Xymon are poller / pinger frameworks that come with built in scripts to verify service availability.
- Snmptrapd running on Linux receives and logs traps to /var/log/messages or whatever
- SNMPTT (SNMP Trap Translator) processes the logs and sends events to either Nagios or Xymon
- http://www.snmptt.org/about.shtml
- http://www.nagios.org/
  http://snmptt.sourceforge.net/docs/snmptt.shtml#Nagios-Netsaint
- http://www.xymon.com/
  http://cerebro.victoriacollege.edu/hobbit-trap.html

OpenNMS

- Open Network Monitor System
- Java based native SNMP monitor
  - Bit of a resource hog, but rich functionality makes up for it
- Native Trap processing and notification functions
- http://opennms.org/

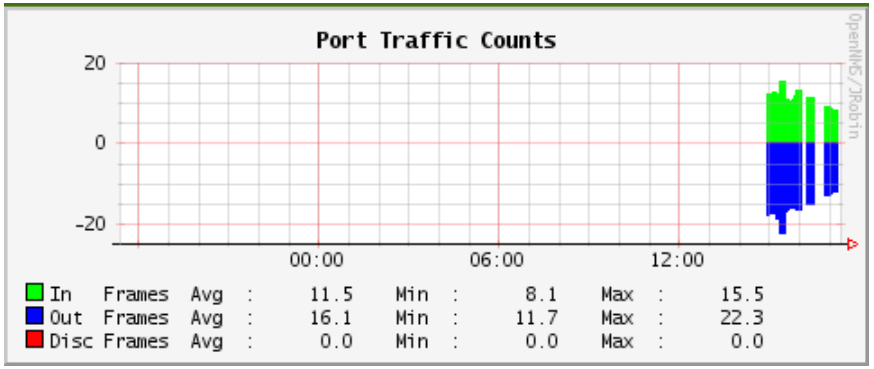

ID	Event ID	Severity	Sent Time	Responder	Respond Time	Node	Interface	Service
1690	662902	Minor	11/6/09 12:42:12 PM			LTICVM9_PRVV71L2	192.168.70.26 [+]	
			LTICVM9_PRVV71L2 snmp link down trap received at Friday, November 6, 2009 12:42:08 PM EST Port number 78 ... I think. Maybe.					
1666	648832	Normal	11/5/09 4:57:15 PM			LTICVM9_PRVV71L2	192.168.70.26 [+]	
			LTICVM9_PRVV71L2 link came up at Thursday, November 5, 2009 4:57:12 PM EST Port number ... I think. Maybe.					

IBM offerings

- Tivoli NetView
  - http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+NetView
- Tivoli Network Manager IP Edition
  - http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+Network+Manager+IP+Edition

ITNM-IP

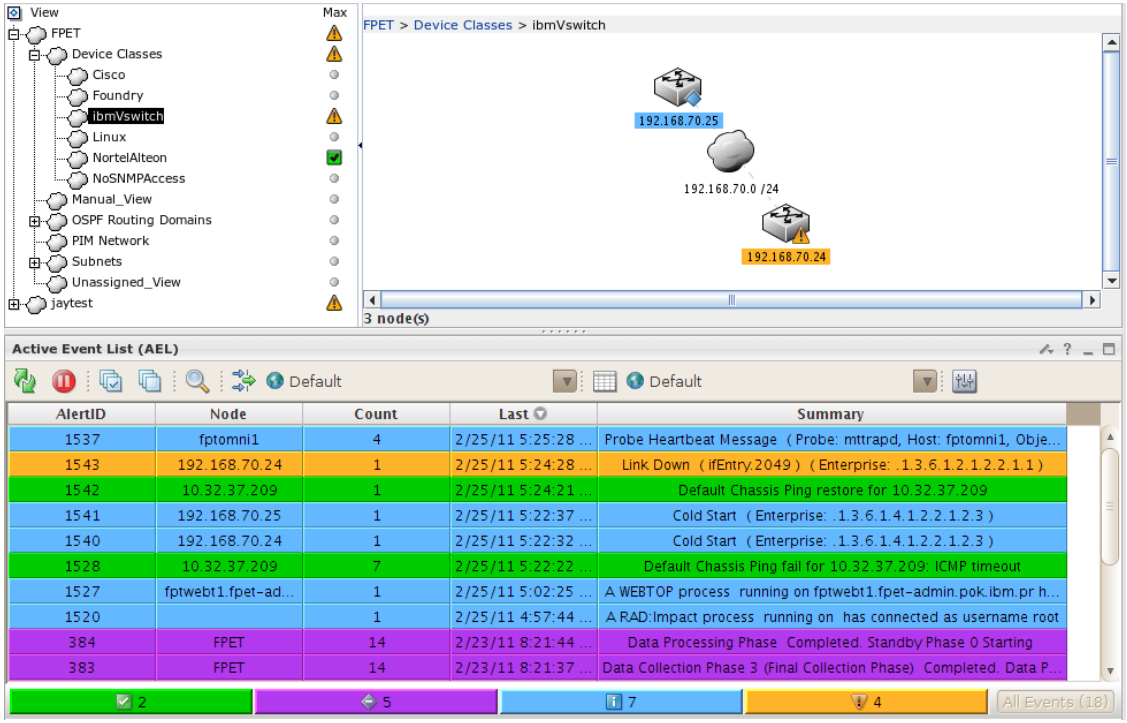

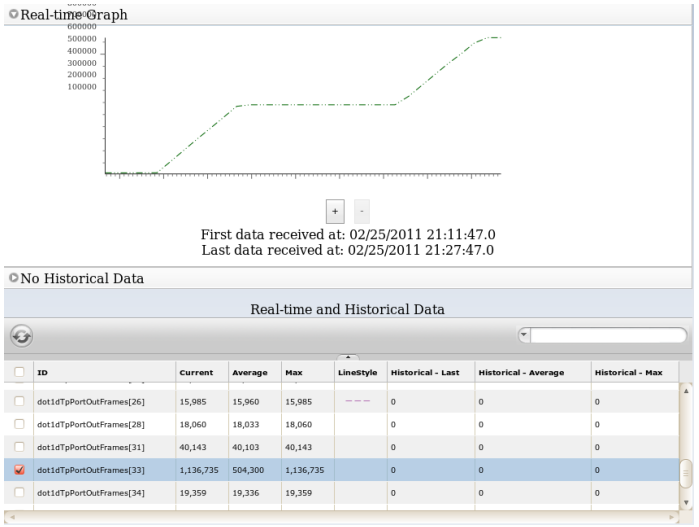

- Integrates with Omnibus for event correlation
- Pulls real time performance counters and compares to historical averages
- Can also generate topology maps for root cause

For more information

- System Z Platform Test library:
  - http://www-03.ibm.com/systems/services/platformtest/servers/systemz_library.html
- The SNMP Paper this presentation is based on:
  - http://www-03.ibm.com/systems/resources/snmp-whitepaper-legal.pdf

This slide intentionally left blank