z/OS IP Load Balancing Advisor (LBA) for External Load Balancers: Background and the SASP Protocol

z/OS Load Balancing Advisor (LBA) - agenda

- **Background**
- **The SASP protocol**

Background

Background information

- **Several existing z/OS IP workload balancing solutions already exist**
  - External load balancing solutions with little or no Sysplex awareness
  - Sysplex-aware external load balancing solutions
    - Such as Network Dispatcher
  - Sysplex-internal load balancing solutions
    - Such as Sysplex Distributor
    - Such as DNS/WLM
- **So why add another?**
  - No one solution is best for everyone - advantages and disadvantages
  - Things to consider...
    - Administration and configuration
    - Support for TCP and UDP applications?
    - Extra network flows?
    - Client/server affinities supported?
    - Network Address Translation required?
    - Use z/OS Network QoS policy?
    - Use WLM (WorkLoad Manager) recommendations?
    - Availability of load balancing when Load Balancing component is unavailable
    - Caching
    - Server health information available?
    - etc.

The rationale behind developing the z/OS Load Balancing Advisor (LBA)

- **Let's assume that you have selected an external IP load balancing solution for your z/OS Sysplex environment**
  - Some possible reasons:
    - Prefer to have a single load balancing solution across multiple platforms in your environment
    - Administration of the load balancing functions belongs to network administration domain (not z/OS administrators)
    - Requirements for content-based load balancing
      - Need to perform load balancing/routing decisions based on data content (inspection of URL, session IDs, cookies, etc.)
      - This is often combined with SSL offloading functions - need to decrypt data prior to inspection
- **But what if the external load balancing solution has no awareness of the Sysplex environment?**
  - Is the Sysplex treated just like any other server cluster?
  - Is it aware of the current/changing workload conditions on the various systems in the Sysplex cluster?
  - Is it aware of the health and status of applications and/or systems?
- **Making the external load balancing solution "Sysplex aware" can help answer many of these questions**
  - The z/OS Load Balancing Advisor is a key component that allows any external load balancing solution to become "Sysplex aware".
    - The external load balancer needs to support the Server Application State Protocol (SASP) to obtain Sysplex information to use in its load balancing decisions.

The SASP protocol

SASP - architecture

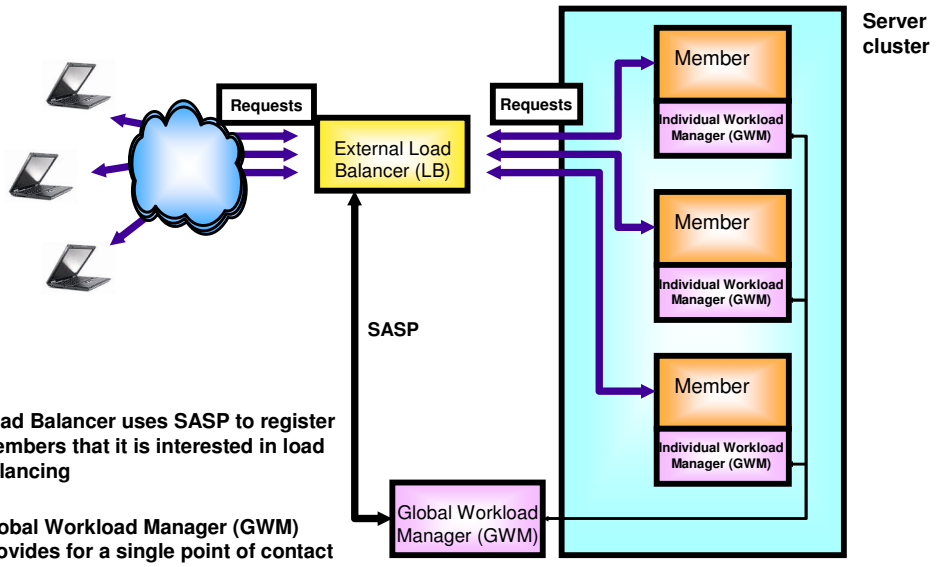

- **Load Balancer uses SASP to register members that it is interested in load balancing**
- **Global Workload Manager (GWM) provides for a single point of contact for the entire cluster**

SASP protocol elements

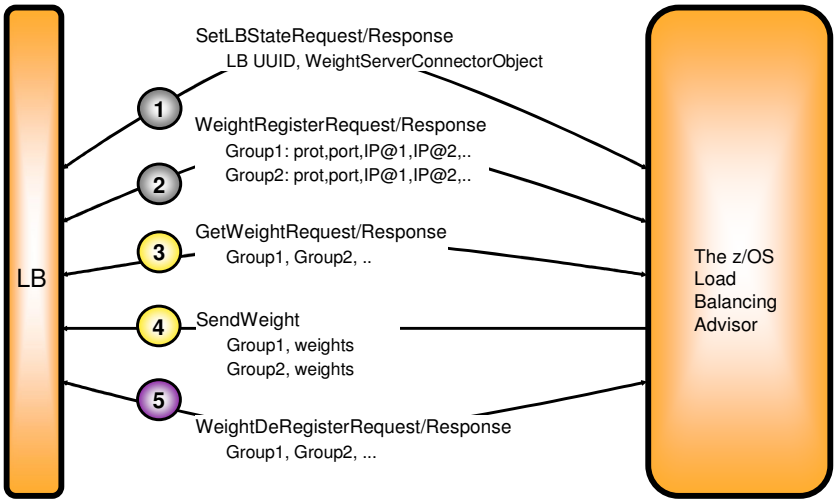

- **1 and 2 performed during LB startup**
- **3 and 4 performed at periodic intervals**
- **5 performed during LB shutdown**

SASP: Server Application State Protocol

Load balancer registrations through SASP

- **The load balancer may register two types of groups for which it wants weights:**
  - A system group
    - Represented by a list of IP addresses only.
    - IP addresses are matched to TCP/IP stacks in the Sysplex.
    - WLM weights for the LPARs are retrieved.
    - CS weight indicates if IP address is active in the Sysplex or not
      - 0 means quiesced
      - 100 means not quiesced
    - LBA displays will show a protocol value of zero for system group registrations.
  - An application group
    - Represented by a list of IP addresses, protocols (TCP or UDP), and ports.
    - Server address spaces are matched to registrations.
    - WLM weights for the LPARs are retrieved.
    - CS weights are calculated factoring in how well the server instances are performing.
    - LBA displays will show protocol as TCP or UDP with the registered port numbers
- **When an external load balancer connects to a global workload manager, it instructs the manager how it wants weights presented:**
  - The load balancer will poll every so often to obtain the current weights
  - The load balancer requests the advisor to push weights down at certain intervals or when the weights change
    - This is how a Cisco CSM external load balancer behaves

SASP update frequency

- **SASP supports both a "push" and a "pull" model for updating the load balancer with workload recommendations**
  - Support of either by the load balancer is implementation dependent
  - Load balancer tells GWM which model it wants to use
  - "Pull" model
    - GWM "suggests" a polling interval to the load balancer
      - z/OS Load Balancing Advisor uses the configurable update_interval value for this purpose
    - Load balancer has the option to ignore this value
    - Load balancer requests updates each polling interval
  - "Push" model
    - GWM sends updated information to the load balancer on an interval basis
      - z/OS Load Balancing Advisor uses the configurable update_interval value for this purpose
    - GWM may send data more frequently than the interval period
  - Load balancer determines whether it wants information about all members it registered or only changed information about its registered members

SASP update frequency *(continued)*

- **Tradeoffs of higher vs. lower update frequency**
  - With more frequent updates...
    - Benefits
      - Workload distribution approaches real-time, reducing latency between when changes appear on the target application and when the load balancer is notified of them
        - Workload distributed more accurately according to system and application capacity
        - Application availability status is more accurate
    - Costs
      - More network flows
        - Between the GWM and the individual workload managers
        - Between the GWM and the load balancers
      - Additional cycles consumed by GWM and LBs
- **Push vs. Pull**
  - May be restricted to one or the other by the LB implementation
  - Pull model may delay when status change information about an application reaches the LB

Products that support the SASP protocol (mid-2005)

- **SASP Global Workload Managers (GWMs)**
  - EWLM (Enterprise Workload Manager)
    - Part of IBM Virtualization Engine 1.0
      - Supported platforms:
        - IBM AIX 5L Version 5.2
        - Microsoft Windows 2000 Advanced Server, 2000 Server, 2003 Enterprise Edition, 2003 Standard Edition
        - Sun Microsystems Solaris 8 (SPARC Platform Edition), 9 (SPARC Platform Edition)
        - Linux on zSeries and System z9 for next release of IBM Virtualization Engine
  - z/OS Load Balancing Advisor
    - Part of z/OS Communications Server (z/OS V1R4 and higher)
- **Load Balancers**
  - Cisco Content Switching Module (CSM) level 4.1 (2.5)
    - Uses the "Push" model
  - Cisco Content Services Switch (CSS)
    - Being tested with z/OS LBA
  - Other vendors likely in the future

EWLM workload balancing

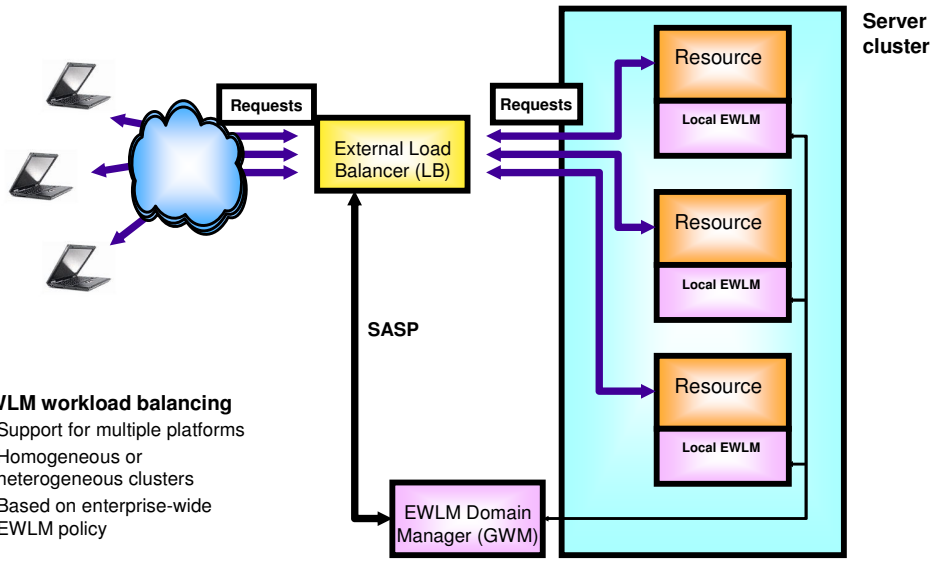

- **EWLM workload balancing**
  - Support for multiple platforms
  - Homogeneous or heterogeneous clusters
  - Based on enterprise-wide EWLM policy

Trademarks, Copyrights and Disclaimers